

Customer-Driven Product Scope

Shresth Garg* Gaurav Khanna† Sagar Saxena‡

March 20, 2026

Click [here](#) for the latest version.

Abstract

We study how buyer-seller relationships shape product introduction decisions by firms. Using administrative records of 54 million transactions from a large Indian state, we document that firms sell new products primarily to existing clients, that product introductions raise sales of complements already in the portfolio, and that firms enter markets where their existing buyers are active. We estimate a structural model in which buyers prefer familiar sellers and value sourcing multiple products from a single supplier. Counterfactual simulations show that product introductions rise when either preference is weakened, and more so in markets where these preferences are strongest.

*University of Pennsylvania, Department of Economics. Contact: shresth@sas.upenn.edu

†UC San Diego, School of Global Policy and Strategy. Contact: gakhanna@ucsd.edu

‡University of Pennsylvania, Wharton School. Contact: sagarsxn@wharton.upenn.edu

1 Introduction

Firms in developing countries stay small. A large body of research attributes this to supply-side constraints: limited access to capital (Banerjee & Duflo, 2014; Buera, Kaboski, & Shin, 2011), factor misallocation (Hsieh & Klenow, 2009), weak management practices (Bloom et al., 2013), and regulatory barriers that penalize expansion (Besley & Burgess, 2004). One path through which firms grow is by expanding their product scope. Multi-product firms account for 39 percent of US manufacturing establishments but 87 percent of output (Bernard, Redding, & Schott, 2010), and here too the existing explanations emphasize supply-side forces: production capabilities (Eckel & Neary, 2010) and input-side forces (Boehm, Dhingra, & Morrow, 2022; Goldberg et al., 2010). Whether demand-side forces also constrain product expansion is less well understood. We ask whether one such force, buyer-seller relationships, shapes which products firms can profitably introduce. If buyers prefer transacting with sellers they already know, and if they value sourcing multiple goods from a single supplier, then a firm’s existing customer relationships and product portfolio define its opportunities for expansion as much as its productive capabilities do.

We study this question using administrative transaction records from India’s electronic waybill (e-way bill) system, which links firms to their specific buyers with rare precision. The system mandates electronic documentation for commercial shipments above a value threshold; each record lists the specific items, quantities, and values in the shipment, along with the identities and locations of both buyer and seller. Our sample covers approximately 54 million shipments over two years from a large Indian state. These records let us identify when firms introduce new products, track whether those products sell to existing clients or new buyers, and observe whether buyers source multiple products from a single seller.

We find three patterns using these data. First, firms sell their new products

primarily to existing clients. In the first month after introduction, 55 percent of sales go to buyers who had previously purchased from the firm, and this share remains above 50 percent after six months. Second, introducing a product increases sales of the firm's other products. Comparing firms that introduce products to those that never do, total firm sales rise by 95 percent around introduction, and three-quarters of this increase comes from products the firm was already selling. These gains are largest for complementary products: those frequently purchased together with the new product see larger sales increases. Third, potential demand from existing clients predicts which products firms introduce. We measure access to buyers for each product and find that firms are roughly fifteen times more likely to introduce a product if their existing clients already buy it elsewhere than if only non-clients do. Together, these patterns point to two demand-side mechanisms: buyers prefer familiar sellers, and buyers value sourcing multiple products from a single supplier.

Next, we estimate a structural model to quantify these mechanisms. The model treats markets as specific product bundles. For example, buyers who need both cement and steel form a distinct market from those who need only cement. In each market, buyers choose a combination of sellers. They can source the entire bundle from one firm or mix and match across multiple suppliers. This structure implies that a firm selling only cement competes in multiple markets simultaneously: it vies for buyers who want only cement, but also for those building a bundle of cement and steel. Introducing steel allows the firm to offer the complete bundle itself, which may be valued by buyers. Two parameters in the buyer demand model capture the mechanisms of interest: a repeat-seller bonus, which captures the value of transacting with existing sellers, and a single-seller bonus, which captures the value that buyers place on sourcing multiple products from the same firm. Buyers weigh these factors against price and distance. On the supply side, firms engage in Bertrand-Nash price competition and decide whether to introduce new products by comparing expected profit gains to a fixed cost. Expected profit gain has two components: direct sales of

the new product and indirect effects on existing products.

We estimate the model in two stages. The first stage recovers demand parameters using a nested fixed-point procedure that matches observed market shares and maximizes the likelihood of observed purchase decisions; we then invert firms' first-order conditions to recover marginal costs. Our estimates show that both the repeat-seller and single-seller bonuses are positive: buyers prefer existing sellers and value bundling from a single source. The second stage estimates the fixed cost distribution from observed entry decisions. For each potential product introduction, we simulate expected profit gain by drawing quality and cost realizations and solving for post-entry equilibrium prices. We then recover the distribution of fixed costs by matching the entry probabilities implied by these profit gains to the actual entry rates in the data.

We use the estimated model to run counterfactual experiments which examine how product introduction would change if buyers cared less about purchasing from existing sellers or about sourcing multiple products from the same firm. In the first experiment, we reduce the repeat-seller bonus by 25 percent, representing an environment where existing relationships confer a smaller advantage. This raises the mean product introduction probability by 50 percent. Incumbents have been around longer and have larger existing customer bases than potential entrants. Reducing the repeat-seller bonus lowers the return on these established relationships. This makes it easier for entrants to compete for buyers and increases the expected profits of entry. The second experiment reduces the single-seller bonus by 50 percent. This also increases aggregate entry, though the effect is smaller. However, the probability of entry rises significantly in markets where incumbents derive a greater share of their revenue from multi-product bundles where the firm supplies all products. In these markets, reducing the bonus erodes the portfolio advantage of multi-product firms and increases the incentives for new firms to enter.

These counterfactual simulations show that buyer preferences for existing sell-

ers and single-sourcing favor established firms by rewarding incumbents with deep customer networks and large portfolios. In contrast, entrants and smaller firms are constrained in their ability to scale. This creates a feedback loop where past success in building a customer base lowers the cost of future expansion. Consequently, market structure may remain concentrated not because incumbents are more productive, but because their demand-side assets (their relationships and portfolios) deter entry.

Related Literature. This paper contributes to work on demand-side constraints on firm growth in developing countries. Goldberg and Reed (2023) argue that development requires adoption of increasing-returns-to-scale technologies, which in turn requires sufficient market size to justify the fixed costs of switching; when domestic and international markets are too small, supply-side investments alone cannot sustain growth. Other work has provided direct evidence: limited market size restricts entry and sustains high markups (Leone, Macchiavello, & Reed, 2025), demand shocks raise profits for firms in crowded sectors (Hardy & Kagy, 2020), and access to foreign buyers raises quality and profitability for developing-country producers (Atkin, Khandelwal, & Osman, 2017). We identify a demand-side mechanism operating through product scope: buyer-seller relationships constrain which goods firms can profitably introduce. A firm capable of making a new product may not enter if its existing customers do not want that product (limiting the returns to bundling), or if potential buyers are tied to other sellers (raising the cost of acquiring new demand).

The mechanism we study is rooted in the literature on buyer-seller relationships and customer capital. Foster, Haltiwanger, and Syverson (2016) show that new plants grow slowly because they must accumulate demand over time; physical productivity alone does not translate into revenue. Gourio and Rudanko (2014) formalize this in a model where firms invest in a customer base that shapes firm dynamics. Switching costs reinforce these patterns: Monarch (2022) estimates that changing suppliers in US-China trade is costly, generating rents for incumbents, and Eaton et al. (2025)

model how search and learning generate the dynamics of export relationships. In developing countries, relational contracting plays a central role (see Macchiavello, 2022, for a survey). McMillan and Woodruff (1999) document that interfirm relationships in Vietnam substitute for weak legal institutions, Macchiavello and Morjaria (2015) show that Kenyan rose exporters prioritize relational buyers over spot-market sales during a supply shock, revealing the value of established relationships, and Startz (2024) shows that search and contracting costs shape trade patterns in Nigeria. This literature studies how relationships affect trade volumes, the terms of exchange, and which transactions take place. We show that these relationships also shape the firm’s product scope: buyer-seller ties determine which products a firm introduces, not just the volume or terms of existing trade.

Finally, our paper contributes to research on multi-product firms. Much of this literature emphasizes supply-side determinants of product range (Bernard, Redding, & Schott, 2010; Boehm, Dhingra, & Morrow, 2022; Eckel & Neary, 2010; Goldberg et al., 2010). On the demand side, Bernard et al. (2019) document “carry-along trade”: manufacturers frequently export products they do not produce, a pattern consistent with demand complementarities favoring multi-product sellers. We build on this by estimating a structural model incorporating both bundling preferences and trading relationships. The model allows us to decompose the forces driving product introduction and to quantify, through counterfactual experiments, how much each demand-side channel contributes to entry barriers.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and data. Section 3 presents descriptive evidence on the role of existing clients and complementarity in product introduction. Section 4 develops the structural model. Section 5 describes the estimation procedure and presents results. Section 6 conducts counterfactual experiments. Section 7 concludes.

2 Setting and Data

To study how firms expand their product portfolios, we use administrative transaction records from India’s e-way bill system. These data track the universe of formal sector shipments above a value threshold, recording the identities and locations of buyers and sellers, the products shipped, and the transaction value. We observe each firm’s complete shipment history and can therefore identify when firms begin selling new products and to whom.

2.1 The E-Way Bill System

In 2017, India introduced a Goods and Services Tax (GST) administered through a centralized digital platform. The reform mandated electronic waybills (e-way bills) for commercial shipments, creating detailed real-time records of goods movements. Any shipment above Rs. 50,000 (approximately \$600) requires an electronic document recording buyer, seller, product, and value. These records form the basis of our data.

Three features of this system generate high-quality data. First, the legal mandate ensures that all formal sector shipments above the threshold are documented. The Rs. 50,000 cutoff is low enough to capture the bulk of business-to-business transactions while filtering out small retail purchases. Second, penalties and physical enforcement at highway checkpoints create compliance incentives: officers can intercept vehicles and detain goods lacking valid documentation. Third, the tax structure makes the system self-policing. Buyers can claim credits for taxes paid on their purchases, but only with valid documentation. This creates pressure on sellers to report accurately.

2.2 Data and Sample

We obtain the universe of e-way bills involving sellers or buyers in a large Indian state for the period April 2018 through March 2020.¹ Each e-way bill records the tax identifiers and locations of both buyer and seller, the value and quantity of goods, and the HSN (Harmonized System of Nomenclature) code classifying the product.

We define firms at the establishment-level; in our sample, a firm is a combination of tax identifier (GSTIN) and location identifier (PIN code). The same company operating from two locations appears as two distinct firms. This approach treats each establishment as a separate economic unit, which aligns with our interest in understanding location-specific product decisions.

2.3 Defining Product Introductions

A product introduction is a firm’s first shipment of a product category that was not part of its initial portfolio. To operationalize this concept, we impose three conditions. First, the introduction must occur at least six months after the firm’s initial appearance in the data. This lag ensures we observe expansion rather than the firm’s starting product mix. Second, the firm must ship the product in more than half of the months following introduction. This requirement distinguishes genuine additions from experimental shipments that the firm quickly abandoned. Third, the introduction must happen at least six months before the end of our sample, so we have adequate time to observe post-introduction outcomes. We provide details on how we construct our sample for reduced-form analysis and structural estimation in [Appendix A](#).

¹We end the sample in March 2020 to avoid the disruptions of the COVID-19 pandemic. See [Appendix A](#) for details on raw data characteristics and sample construction.

3 Descriptive Evidence

Using this sample of product introductions, we explore two questions. First, what happens after introduction? Does the firm rely on existing clients to purchase the new product, and does the introduction affect sales of existing products? Second, what predicts which products firms introduce? Do geographic proximity, client relationships, and complementarity shape these decisions? To guide our empirical analysis, in Appendix B, we develop a stylized model that formalizes these intuitions and generates testable predictions.

We find three patterns. Firms rely heavily on existing clients when selling new products, with over half of initial sales going to buyers with prior relationships. Product introductions increase sales of complementary products already in the firm’s portfolio. And market access to existing clients, particularly for complementary products, predicts which products firms introduce. These patterns inform the structural model we develop in the next section.


3.1 Sales following new product introduction

We denote sellers by i , products by p , and time by t . The month of a product introduction is normalized to $t = 0$.

3.1.1 Client Composition

Product introduction differs from entry into a new market because introducing firms have an existing customer base. We begin by quantifying how important these relationships are. If existing clients derive utility from purchasing from familiar sellers,

Figure 1: *Share of Sales to Existing Clients*



Notes. This figure plots the share of new product sales going to buyers who previously purchased from the firm. Month 0 marks product introduction; pre-introduction shares are set to zero by construction. 95% bootstrap confidence intervals are shown.

they should account for a large share of new product sales.

We define existing clients as those who purchased from the firm in any month prior to the product introduction. The share of sales to existing clients, for product p by firm i in month t , is the fraction of that product's monthly sales going to buyers who previously purchased from the firm. We set this share to zero for pre-event periods ($t < 0$). Figure 1 plots the average of this share across all product introductions. Firms rely heavily on their existing customer base when launching new products. In the first month after introduction, 55% of sales go to existing clients. This share declines only modestly (to roughly 50%) over the following six months, suggesting that existing relationships remain an important source of demand well beyond the initial launch.

3.1.2 Sales of existing products

Existing clients purchase new products, but does product introduction affect the firm’s existing business? If buyers prefer to source bundles from a single seller, introducing a complementary product should increase sales of existing products.

We consider two outcome variables: *overall sales*, the sum of sales across all products in a given firm-month, and *prior sales*, sales of products that the firm was selling prior to introduction. We use an event study design to estimate how outcomes evolve around product introduction. For firm i that introduces a product, we estimate:

$$y_{i,t} = \alpha_i + \alpha_t + \sum_{\tau=-k}^k \beta_{\tau} D_{i,t}^{\tau} + \varepsilon_{i,t} \quad (1)$$

where $D_{i,t}^{\tau} = \mathbb{1}[t - G_i = \tau]$ and G_i is the month of product introduction. We use the Callaway and Sant’Anna (2021) estimator, comparing treated firms to control firms that never introduce new products during our sample period. We transform both outcome variables using the inverse hyperbolic sine.

Figure 2 presents the results. Product introduction increases overall sales by 95% and prior sales by 75%. The gap between these two figures implies that roughly 20 percentage points come from the new product itself, while 75 percentage points represent spillovers to existing products. Most of the sales increase comes from spillovers to existing products rather than from the new product itself.

The theory predicts that this increase in prior sales should concentrate among complements. To test this, we need a measure of product complementarity. We use our transaction data to identify which products are purchased together in practice. Specifically, we measure complementarity using lift scores, which capture how much more likely two products are to be purchased together than would be expected if

purchases were independent. For products p and p' , the lift score is:

$$\text{lift}(p, p') = \frac{\Pr(p \cap p')}{\Pr(p) \times \Pr(p')} \quad (2)$$

where $\Pr(p)$ denotes the probability that product p is purchased in a given transaction, and $\Pr(p \cap p')$ denotes the probability that both products p and p' are purchased together in the same transaction. Higher lift scores indicate stronger complementarity.

To avoid mechanical relationships, we compute lift scores excluding all shipments from firms that introduce products during our sample period.² For each new product, we use the 90th percentile of its lift distribution as the threshold to classify other products as high complementarity (above threshold) or low complementarity (below threshold).³

We run separate event studies for sales of high and low complementarity products, restricting to treated firms that sell both types. Figure 2 shows that sales of both types increase, but the increase is substantially larger for high complementarity products. This differential effect by complementarity is consistent with bundled purchasing: buyers who value sourcing multiple products from a single seller increase their purchases when the firm adds a complement.

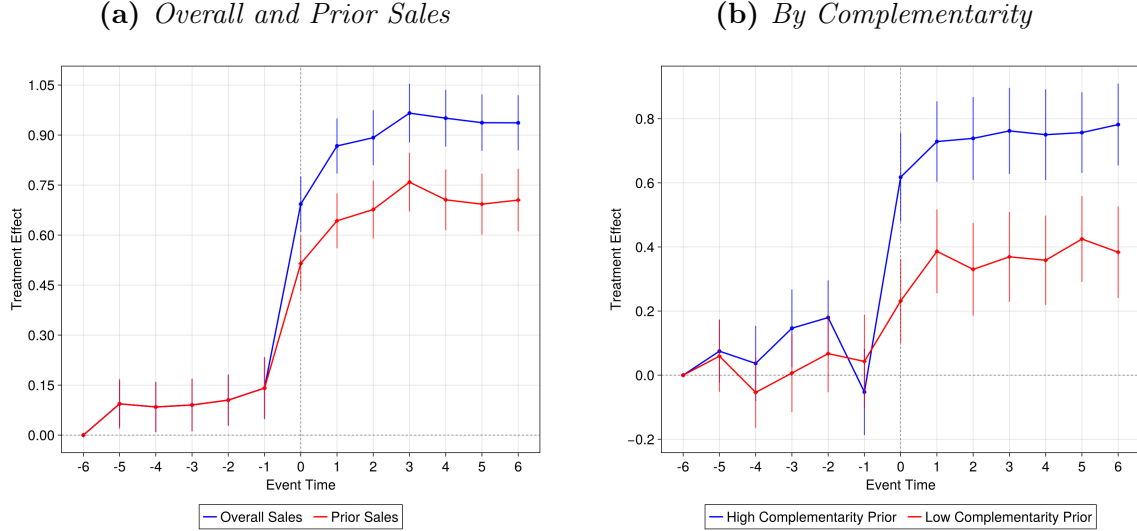
3.1.3 Prices

Having established that product introduction increases sales of existing products, we now ask whether firms adjust prices. Bundling creates opposing incentives: lowering

²If we included these firms, a mechanical bias would arise: firms that introduce a new product and sell it alongside existing products would generate co-purchase patterns reflecting their product portfolio decisions rather than underlying demand complementarity. This would inflate lift scores for product pairs sold by introducing firms, biasing our complementarity measure toward finding stronger effects precisely for the products we study.

³A significant proportion of lift scores are zero because many product pairs are never purchased together. In cases where the 90th percentile of lift scores for product p is zero, we classify all products with positive lift scores as having high complementarity with p .

Figure 2: *Treatment Effects on Sales Around New Product Introduction*




Notes. Panel (a) shows treatment effects on overall sales and sales of products the firm sold prior to introduction. Panel (b) decomposes prior sales by complementarity with the newly introduced product, where high complementarity denotes products with lift scores above the 90th percentile. Both panels report Callaway and Sant’Anna (2021) estimates with month 0 marking introduction. All outcomes are transformed using the inverse hyperbolic sine.

prices attracts more buyers to the bundle, but raising prices extracts surplus from buyers who value bundle convenience. The net effect is an empirical question.

We define *bundled price* as the price of an existing product when sold together with the newly introduced product in the same transaction. We compare bundled prices to prices of the same product when sold separately, isolating the effect of bundling on prices while holding product characteristics fixed. The sample includes all firm-products with at least one sale every month, whether bundled or unbundled.

Figure 3 presents the event study results. Prices of existing products decrease when bundled with the newly introduced product. This suggests that the demand expansion force dominates: firms lower prices on existing products to increase sales of both products via the bundle.

Figure 3: *Treatment Effects on Bundling Prices*



Notes. This figure shows how prices of existing products change when sold together with the newly introduced product in the same transaction, relative to sales of those products sold separately. The sample includes firm-products with at least one observation in every month of the event window. Estimates use the Callaway and Sant’Anna (2021) estimator.

3.2 Market Access and Product Introduction

The previous results document what happens after introduction. We now ask whether firms anticipate these benefits when deciding which products to add. If firms understand that existing clients and complementarity raise returns to product introduction, these factors should predict introduction decisions.

Testing this requires measuring each firm’s access to relevant demand. A firm’s potential buyers are scattered geographically, and closer buyers are presumably easier to serve. We therefore construct a gravity-based measure, *market access*, that sums

over buyers weighted by their purchasing volume and discounted by distance. A firm with high market access faces strong nearby demand; a firm with low market access faces either weak demand or distant buyers. For seller i and product p , we define:

$$\text{Market Access}_{i,p} = \sum_{j \in \mathcal{J}} \frac{\text{asinh}(\text{purchases}_{j,p})}{(\text{distance}_{ij})^2} \quad (3)$$

where asinh is the inverse hyperbolic sine, $\text{purchases}_{j,p}$ is the average monthly purchase of product p by buyer j , and distance_{ij} is the geographic distance between seller i and buyer j .

We compute market access using transaction data from each firm’s first six months in the sample. This window precedes any product introduction, since our treatment definition requires introductions to occur at least six months after firm entry. Existing clients are buyers who purchased from the firm during this initial period; all other buyers are classified as new.⁴ The set \mathcal{J} varies by specification: all buyers, existing clients only, or new buyers only.

Each firm can potentially introduce thousands of products, so expanding the choice set generates a very large number of firm-product observations. To make estimation tractable, we retain all treated firms but randomly subsample control firms, keeping four control firms for each treated firm. We apply inverse probability weights to recover population-level estimates.

We estimate:

$$\text{introduced}_{i,p} = \beta \cdot \mathbf{1}[\text{High Market Access}_{i,p}] + \gamma \mathbf{X}_{i,p} + \lambda + \varepsilon_{i,p} \quad (4)$$

where $\text{introduced}_{i,p}$ is an indicator for whether seller i introduces product p . High market access is defined as above the median within each product category, so the

⁴This definition differs from the event study, where existing clients are defined relative to the introduction date. We cannot use that definition here because control firms—those that never introduce products—have no introduction date. Using the first six months after entry provides a consistent baseline across treated and control firms. This choice is conservative: clients acquired between month six and the introduction date are classified as new buyers, understating the role of existing relationships.

Table 1: *Market Access and Product Introduction*

	Introduced Product (Basis Points)			
	(1)	(2)	(3)	(4)
High Market Access (Product)	0.33394*** (0.055)	0.26880*** (0.044)	0.30656*** (0.041)	0.33847*** (0.042)
High Market Access (All Products)		0.09067 (0.050)		
Product (HSN4) Fixed Effects	Yes	Yes	Yes	Yes
Pincode Fixed Effects			Yes	Yes
Firm Fixed Effects				Yes
N	4,038,457	4,038,457	4,038,457	4,038,457
R^2	0.000	0.000	0.002	0.006
Mean of Dep. Var.	0.506	0.506	0.506	0.506

Notes. The dependent variable is an indicator for whether the firm introduces the product; coefficients represent changes in introduction probability in basis points. High market access is defined as above the median within each product category. Column 2 includes market access to all products as a control. Standard errors clustered by firm in parentheses. Inverse probability weights account for subsampling of control firms. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

comparison is between firms with high versus low market access for the same product. We progressively add fixed effects to isolate different sources of variation.

Table 1 shows that access to potential demand is positively associated with product introduction. All columns include product fixed effects, which compare firms with high versus low market access for the same product. The first column shows that high market access increases the probability of product introduction by approximately 0.3 basis points. With a baseline introduction probability of 0.5 basis points, this represents a 60% increase. The second column adds market access to all products

as a control, capturing total potential demand in the firm’s location; this coefficient is small and insignificant, indicating that product-specific demand potential, rather than simply being located in a high-demand area, drives introduction decisions. The third column adds pincode fixed effects to absorb all location-specific differences and compare firms in the same location considering the same product. The coefficient remains statistically significant. The fourth column adds firm fixed effects, comparing products within the same firm. This is our preferred specification, with a coefficient of 0.34 basis points.

Table 2 decomposes market access by buyer type. We construct the same measure as before, but now separately summing over existing clients and new buyers. The first three columns progressively add product, pincode, and firm fixed effects. In our preferred specification with all three fixed effects, the coefficient on existing clients is 2.9 basis points, while the coefficient on new buyers is 0.2 basis points. Thus, firms appear to weight existing relationships heavily when choosing products.

The final two columns add product-by-pincode fixed effects, comparing firms in the same location considering the same product. Since market access is constructed from distance to buyers, firms in the same pincode have identical distances to all potential buyers. This means there is little variation in market access to new buyers across firms in the same location—they all face similar pools of potential new buyers. In contrast, existing client relationships vary substantially across firms even within the same location. The coefficient on existing clients remains large and precisely estimated (3.2–3.5 basis points).

Table 3 tests the role of complementarity in product introduction. We use the same lift-based measure of complementarity from the previous analysis looking at impact on sales. For each product p , we compute lift scores with all other products and classify those above the 90th percentile as high complementarity. We then construct market access to high complementarity products by summing over buyers

Table 2: Market Access: Existing Clients vs New Buyers

	Introduced Product (Basis Points)				
	(1)	(2)	(3)	(4)	(5)
High Market Access (Existing Clients)	2.62296*** (0.199)	2.80017*** (0.208)	2.88213*** (0.212)	3.21959*** (0.312)	3.50887*** (0.366)
High Market Access (New Clients)	0.09696 (0.062)	0.17489*** (0.042)	0.21054*** (0.043)	-0.78079** (0.254)	-0.43947 (0.288)
Product (HSN4) Fixed Effects	Yes	Yes	Yes	Yes	Yes
Pincode Fixed Effects		Yes	Yes	Yes	Yes
Firm Fixed Effects			Yes		Yes
Product (HSN4) × Pincode Fixed Effects				Yes	Yes
<i>N</i>	4,038,457	4,038,457	4,038,457	3,155,505	4,038,457
<i>R</i> ²	0.000	0.002	0.006	0.105	0.236
Mean of Dep. Var.	0.506	0.506	0.506	0.506	0.506

Notes. The dependent variable is an indicator for whether the firm introduces the product; coefficients represent changes in introduction probability in basis points. Market access is computed separately for existing clients (buyers who purchased during the firm's first six months) and new buyers. High market access is defined as above the median within each product category. Standard errors clustered by firm in parentheses. Inverse probability weights account for subsampling of control firms. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Market Access with Complementarity

	Introduced Product (Basis Points)		
	(1)	(2)	(3)
High Market Access (Product)	0.25580*** (0.047)	0.32421*** (0.046)	0.35993*** (0.047)
High Market Access (High Complementarity)	0.15252** (0.053)	0.24465** (0.088)	0.30798*** (0.084)
Product (HSN4) Fixed Effects	Yes	Yes	Yes
Pincode Fixed Effects		Yes	Yes
Firm Fixed Effects			Yes
N	3,603,740	3,603,740	3,603,740
R^2	0.000	0.002	0.006
Mean of Dep. Var.	0.506	0.506	0.506

Notes. The dependent variable is an indicator for whether the firm introduces the product; coefficients represent changes in introduction probability in basis points. High complementarity products are those with lift scores above the 90th percentile relative to the focal product. High market access is defined as above the median within each product category. Standard errors clustered by firm in parentheses. Inverse probability weights account for subsampling of control firms. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

who purchase products with high lift relative to p . The first column includes product fixed effects, while the second and third columns progressively add pincode and firm fixed effects. In our preferred specification with all three fixed effects, the coefficient on high complementarity market access is 0.31 basis points and the coefficient on product-specific market access is 0.36 basis points; both are statistically significant at the 1% level. Relative to the baseline introduction probability of 0.5 basis points, these represent increases of roughly 60% and 70%, respectively. Thus, complementarity provides additional explanatory power beyond access to buyers of the focal

product.

Table 4 interacts the buyer-type and complementarity decompositions. We construct four measures: market access to existing clients and new buyers for the focal product (as in Table 2), and market access to existing clients and new buyers for high-complementarity products. The latter two measures capture access to buyers of complementary products, further decomposed by whether those buyers have an existing relationship with the firm. The first three columns progressively add product, pincodes, and firm fixed effects. In our preferred specification with all three fixed effects, the coefficient on existing clients is 2.9 basis points and the coefficient on high-complementarity existing clients is 0.51 basis points.

The final two columns add product-by-pincodes fixed effects, comparing firms in the same location considering the same product. Most of the variation in complementarity market access comes from existing clients. In our preferred specification with all fixed effects, the coefficient on existing clients is 3.6 basis points and the coefficient on high-complementarity existing clients is 0.94 basis points.

3.3 Summary

These results establish three patterns. First, firms rely on existing clients for over half of new product sales, suggesting that relationship-specific advantages shape returns to product introduction. Second, introducing a product increases sales of complementary products, so firms cannot evaluate new products in isolation. Third, market access to existing clients, particularly for complementary products, predicts which products firms choose to introduce.

These patterns can be rationalized by the following mechanisms: buyers prefer to purchase from familiar sellers and to consolidate purchases with a single supplier.

Table 4: Market Access: Buyer Type and Complementarity

	Introduced Product (Basis Points)				
	(1)	(2)	(3)	(4)	(5)
High Market Access (Existing Clients)	2.59055*** (0.207)	2.77404*** (0.211)	2.91994*** (0.217)	3.14266*** (0.302)	3.57677*** (0.372)
High Market Access (New Clients)	0.06523 (0.048)	0.18781*** (0.047)	0.22029*** (0.048)	-0.85230** (0.288)	-0.55209 (0.331)
High Market Access (High Comp., Existing Clients)	0.14901* (0.069)	0.35825*** (0.082)	0.51229*** (0.065)	0.38331** (0.129)	0.94372*** (0.176)
High Market Access (High Comp., New Clients)	-0.02849 (0.059)	0.16558 (0.088)	0.25010** (0.083)	-0.35830 (0.415)	0.02898 (0.392)
Product (HSN4) Fixed Effects	Yes	Yes	Yes	Yes	Yes
Pincode Fixed Effects		Yes	Yes	Yes	Yes
Firm Fixed Effects			Yes		Yes
Product (HSN4) \times Pincode Fixed Effects				Yes	Yes
N	3,603,740	3,603,740	3,603,740	2,818,733	3,603,740
R^2	0.000	0.002	0.007	0.104	0.236
Mean of Dep. Var.	0.506	0.506	0.506	0.506	0.506

Notes. The dependent variable is an indicator for whether the firm introduces the product; coefficients represent changes in introduction probability in basis points. Market access is computed separately by buyer type (existing clients vs new buyers) and by complementarity (focal product vs high-complementarity products). High market access is defined as above the median within each product category. Standard errors clustered by firm in parentheses. Inverse probability weights account for subsampling of control firms. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The next section develops a structural model that formalizes these mechanisms.

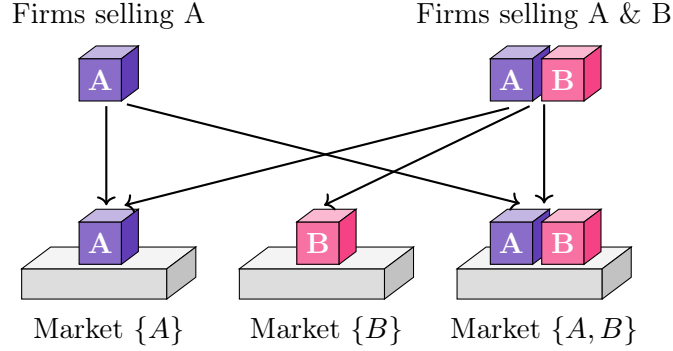
4 Model

We develop a static model of product introduction that captures three features documented in the previous section: the advantage existing clients confer for new product sales, the demand complementarity between products in a firm’s portfolio, and the role of geographic proximity in shaping which products firms introduce. The model incorporates these mechanisms through a relationship bonus β_r and a single-seller bonus γ . On the demand side, buyers purchase bundles of goods and choose which sellers to source from. These choices depend on geographic distance to sellers, existing client relationships, and a preference for sourcing the entire bundle from a single firm. On the supply side, firms engage in Bertrand-Nash price competition and decide which new varieties to introduce after weighing expected variable profits against fixed costs of introduction.

4.1 Setup

We denote HSN codes by h , firms by f , and time periods by t . A *good* is an HSN code while a *variety* is a firm-good pair (f, h) , representing the particular offering that firm f provides for good h . Buyers may require a single good or multiple goods at once. A *bundle* $b = (h_1, \dots, h_K)$ is an ordered collection of one or more goods that a buyer requires. Single-good bundles correspond to standalone purchases; multi-good bundles correspond to buyers who need multiple different goods in a single transaction. For each good in the bundle, the buyer must choose which firm’s variety to purchase.

Figure 4: *Markets Defined by Product Bundles*



Notes. Firms compete in all markets where their products are demanded. Single-product firms selling only A compete in markets $\{A\}$ and $\{A, B\}$. Multi-product firms selling both A and B compete in all three markets.

A *market* (b, t) is defined by a bundle b and a time period t . Figure 4 illustrates the market structure. Because markets are defined by bundles, the same product can appear in multiple markets. For example, a product purchased both alone and as part of a bundle participates in two distinct markets. Each market has a set of buyers $i \in I_b$, and each buyer demands specific quantities $\{q_{i,h_k}\}_{k=1}^K$ of the goods in the bundle. We take quantities as given; the buyer's decision is which variety to purchase for each good.

The model proceeds in two stages. In the first stage, firms decide which new goods to add to their portfolios. Each firm currently sells a set of goods \mathcal{H}_f , offering one variety per good. Firms consider expanding into goods they do not currently sell; adding good h means creating a new variety (f, h) . In the second stage, buyers observe which varieties are available (including any new varieties that firms chose to introduce) and make purchase decisions. We begin with the second stage, characterizing demand and pricing for a given set of varieties, then turn to the entry decision that determines which varieties are offered.

4.2 Demand

We focus on a single market (b, t) and suppress b and t from notation where clear. For each good in a bundle, buyers choose which firm's variety to purchase. When a bundle contains multiple goods, buyers may source all goods from a single firm or divide purchases across multiple firms. We represent this choice as a *seller combination* $\mathbf{f} = (f_1, \dots, f_K)$, where f_k is the firm supplying good h_k . The choice set \mathcal{C}_b consists of all feasible seller combinations—those where each firm f_k sells good h_k .

Buyer i 's utility from seller combination \mathbf{f} depends on the varieties purchased, the distance to each seller, whether the buyer has transacted with each seller before, and whether all goods come from the same firm. Since buyers demand different quantities of each good, we weight each good by its share of total quantity: $w_{ik} = q_{i,h_k} / \sum_{k'} q_{i,h_{k'}}$ is the quantity weight of good k in buyer i 's bundle. Formally, $U_{i\mathbf{f}} = V_{i\mathbf{f}} + \varepsilon_{i\mathbf{f}}$, where V is the deterministic component and ε is a taste shock. The deterministic component is:

$$V_{i\mathbf{f}} = \sum_{k=1}^K w_{ik} \cdot v_{i,f_k,h_k} + \gamma \cdot \mathbf{1}[\text{single}] + \beta_d \cdot \log \left(\sum_{k=1}^K w_{ik} \cdot d_{i,f_k} \right) \quad (5)$$

$$v_{i,f_k,h_k} = \delta_{f_k,h_k} + \beta_r \cdot \mathbf{1}[\text{repeat}_{i,f_k}]$$

$$\delta_{fh} = \alpha \cdot p_{fh} + \xi_{fh}$$

The first equation has three terms. The first term aggregates per-good utilities v_{i,f_k,h_k} across goods in the bundle, weighted by quantity shares w_{ik} . Per-good utility v_{i,f_k,h_k} is buyer i 's utility from purchasing good h_k from firm f_k . It consists of two components: the mean utility δ_{f_k,h_k} of variety (f_k, h_k) , and a relationship bonus $\beta_r \cdot \mathbf{1}[\text{repeat}_{i,f_k}]$ that adds utility when buyer i has previously purchased from firm f_k . Mean utility δ_{fh} decomposes into price p_{fh} and unobserved variety quality ξ_{fh} .

The second term $\gamma \cdot \mathbf{1}[\text{single}]$ is a single-seller bonus that adds utility when all goods in the bundle come from the same firm. The third term $\beta_d \cdot \log(\sum_k w_{ik} \cdot d_{i,f_k})$

captures weighted distance to sellers, where d_{i,f_k} is the distance between buyer i and firm f_k .

We assume taste shocks $\varepsilon_{i\mathbf{f}}$ are independent and identically distributed Type-I extreme value across buyers and seller combinations. The probability that buyer i chooses seller combination \mathbf{f} is:

$$s_{i\mathbf{f}} = \frac{\exp(V_{i\mathbf{f}})}{\sum_{\mathbf{f}' \in \mathcal{C}_b} \exp(V_{i\mathbf{f}'})} \quad (6)$$

For estimation and counterfactuals, we also require the marginal probability that buyer i purchases good h from firm f . This marginal probability sums over all seller combinations where f supplies good h :

$$S_{ifh} = \sum_{\mathbf{f}' \in \mathcal{C}_b} s_{i\mathbf{f}'} \cdot \mathbf{1}[f'_h = f] \quad (7)$$

where f'_h denotes the firm supplying good h in seller combination \mathbf{f}' .

Two features of demand are central to our analysis. First, existing client relationships provide a demand-side advantage. Buyers who have previously transacted with a firm receive additional utility from any subsequent purchase, captured by the relationship bonus β_r . This creates a built-in customer base for new products: when a firm introduces a new good, its existing clients are predisposed to purchase from the firm rather than from competitors. Second, demand complementarities arise from bundled purchasing. The single-seller bonus γ captures the preference for sourcing multiple goods from the same firm, reflecting reduced transaction costs, simplified logistics, or relationship benefits. Consider a buyer requiring goods A and B . A firm selling only good A competes at a disadvantage against competitors offering both goods. Introducing good B allows the firm to offer this bundle convenience. This creates demand complementarity: adding B to the portfolio not only generates direct sales of B but also increases sales of A by making the firm's bundle more attractive.

4.3 Supply

We now turn to pricing. Firms engage in Bertrand-Nash competition and simultaneously choose prices to maximize profits. Each firm f sets a single price p_{fh} for good h , which applies to all sales regardless of the bundle in which the variety appears.

Firm f 's total profit Π_f aggregates across bundle markets b where the firm sells at least one variety. Within each market, it sums over buyers i and seller combinations \mathbf{f}' , weighting each combination by the probability $s_{i\mathbf{f}'}$ that buyer i chooses it:

$$\Pi_f = \sum_{b \in \mathcal{B}} \sum_{i \in I_b} \sum_{\mathbf{f}' \in \mathcal{C}_b} s_{i\mathbf{f}'} \left(\sum_{k=1}^K \mathbf{1}[f'_k = f] \cdot (p_{fh_k} - mc_{fh_k}) \cdot q_{i,h_k} \right) \quad (8)$$

The inner sum computes the contribution from each good in the bundle. The indicator $\mathbf{1}[f'_k = f]$ equals one when firm f supplies good k in seller combination \mathbf{f}' , so the firm earns a profit only on goods it actually sells. Each good contributes $(p_{fh_k} - mc_{fh_k}) \cdot q_{i,h_k}$: price minus marginal cost mc_{fh} , times the quantity q_{i,h_k} that buyer i demands.

Firms set prices to maximize profit. The first-order conditions yield a system of equations linking prices, marginal costs, and quantities (see Appendix C for the full derivation):

$$\mathbf{\Omega}_f(\mathbf{p}_f - \mathbf{mc}_f) = \mathbf{Q}_f/\alpha \quad (9)$$

where \mathbf{p}_f and \mathbf{mc}_f are the vectors of prices and marginal costs for firm f 's products, \mathbf{Q}_f is the vector of quantities sold, and $\mathbf{\Omega}_f$ is a matrix capturing how prices affect demand within the firm's portfolio.

To illustrate, consider a firm selling two goods, A and B . The first-order condition for good A is:

$$\Omega_{AA}(p_A - mc_A) + \Omega_{AB}(p_B - mc_B) = Q_A/\alpha$$

The first term captures the own-price effect that governs how raising p_A reduces demand for A . The second term is the cross-price effect. It captures the impact of

raising p_A on the demand for B . This effect operates through bundles containing both goods—when buyers shift away from seller combinations that include the firm for A , they also reduce purchases of B . Multi-product firms internalize these cross-price effects when setting prices.

4.4 Entry

Let \mathcal{H}_f denote firm f 's current product portfolio. Introducing a new product $h \notin \mathcal{H}_f$ requires a fixed cost F_{fh} and generates a change in variable profit. This change in profits consists of a direct effect and an indirect effect. The direct effect is profit from selling the new product h itself. The indirect effect operates through bundling complementarities: adding h makes the firm more competitive for buyers who demand bundles containing both h and products already in \mathcal{H}_f . These buyers can now source the entire bundle from firm f , receiving the single-seller bonus γ , which raises demand for the firm's existing products. The indirect effect is unambiguously positive because markets are defined by bundles with fixed buyer populations who demand fixed quantities. Buyers choose only which firm to source from; adding product h never makes firm f less attractive for any product $A \in \mathcal{H}_f$.

There are two key forces in our setting that govern how much profits change when a firm adds a new product. First, existing clients receive the relationship bonus β_r , which raises their purchase probability for the new product. Firms with larger client bases therefore earn higher direct profits from product introductions. Second, if the firm already sells complements, introducing h creates single-seller options for buyers of bundles containing both products. The single-seller bonus γ makes the firm's bundle more attractive relative to competitors who cannot offer the same convenience, amplifying the indirect effect on existing product sales.

Prior to entry, the firm does not observe the quality ξ_{fh} or marginal cost mc_{fh} it

will realize. We assume the firm draws (ξ_{fh}, mc_{fh}) from the empirical distribution G_h of incumbents selling good h . Pre-entry profit $\Pi_f(\mathcal{H}_f)$ aggregates across all bundle markets where the firm currently participates, following Equation (8). Post-entry profit $\Pi_f(\mathcal{H}_f \cup \{h\}; \xi, mc)$ depends on the realized draw and is computed at the Bertrand-Nash equilibrium prices that prevail after entry. Expected profit gain is:

$$E[\Delta\Pi_f(h)] = E_{G_h} [\Pi_f(\mathcal{H}_f \cup \{h\}; \xi, mc)] - \Pi_f(\mathcal{H}_f) \quad (10)$$

Fixed costs are drawn from a lognormal distribution with CDF $\Phi_{FC}(\cdot; \mu, \sigma)$. Firm f introduces good h if the realized fixed cost falls below expected profit gain. The probability of entry is therefore:

$$\Pr(\text{entry}) = \Phi_{FC}(E[\Delta\Pi_f(h)]; \mu, \sigma)$$

Thus, higher expected profits increase the probability of entry. Note that we do not model strategic interactions in entry decisions. Each firm takes competitors' portfolios as fixed, which is appropriate given the low entry rates in our data: simultaneous entry by multiple firms into the same product in the same period is extremely rare.⁵

5 Estimation

We estimate the model in two stages. The first stage estimates demand parameters and recovers marginal costs from observed purchase decisions. The second stage estimates the fixed cost distribution from observed entry decisions.

Estimation requires defining valid bundles and classifying observed purchases. We define valid bundles as product combinations that appear together on a single e-way bill. We then assign each transaction to a bundle. If a single e-way bill contains a valid

⁵Two or more firms introduce a new product in the same market (product x month) in less than 0.14% of markets.

bundle, we classify the purchase as a same-seller bundle. If multiple transactions from different sellers within the same month form a valid bundle, they are matched into a multi-seller bundle. Otherwise, transactions are labeled as single-product purchases.⁶

5.1 Demand and Marginal Cost

This section estimates three sets of parameters: the preference parameters $(\beta_d, \beta_r, \gamma)$ on distance, relationships, and bundling; the price coefficient α ; and marginal costs $\{mc_{fht}\}$.

We estimate these parameters in three steps. The first step recovers preference parameters and mean utilities jointly using a nested fixed-point procedure. For a given candidate $(\beta_d, \beta_r, \gamma)$, we solve for the mean utilities that rationalize observed market shares, then evaluate the likelihood of observed purchase decisions. The second step decomposes the recovered mean utilities into price and unobserved quality components, yielding the price coefficient α . The third step inverts firms' first-order conditions to recover marginal costs.

5.1.1 Step 1: Preference Parameters $(\beta_d, \beta_r, \gamma)$

We estimate $\boldsymbol{\theta} = (\beta_d, \beta_r, \gamma)$ using a nested fixed-point procedure in the spirit of Berry, Levinsohn, and Pakes, 1995. For each candidate $\boldsymbol{\theta}$, an inner loop recovers the mean utilities that rationalize observed market shares; an outer loop then evaluates the likelihood of observed purchase decisions given these mean utilities.

The inner loop recovers mean utilities $\{\delta_{fht}\}$ by matching predicted to observed

⁶The estimation sample uses eight-digit HSN product codes and 12 months of transaction data. We focus on the top 20 single-product and top 5 multi-product bundles by transaction count. We restrict the analysis to bundles of at most two products. When a transaction could match multiple bundles, we prioritize by temporal proximity and assign each transaction to at most one bundle.

market shares. In standard demand estimation, shares are defined within a single market. Here, however, a variety (f, h) appears in multiple bundle markets—both single-product bundles $\{h\}$ and multi-product bundles containing h . We therefore aggregate across all bundle markets containing product h to compute each variety’s share.

Formally, let s_{fht}^{pred} denote the predicted market share of variety (f, h) in period t . This share sums over all bundle markets containing h , weighting each buyer’s marginal purchase probability by the quantity that buyer demands:

$$s_{fht}^{\text{pred}} = \frac{\sum_{b \ni h} \sum_{i \in I_{bt}} S_{ifh}(\boldsymbol{\theta}, \boldsymbol{\delta}) \cdot q_{ih}}{\sum_{b \ni h} \sum_{i \in I_{bt}} q_{ih}}$$

The outer sum is over all bundles b containing product h ; the inner sum is over all buyers i in each bundle market. For each buyer, S_{ifh} is the marginal probability of purchasing product h from firm f (Equation 7), which depends on the choice set available in that buyer’s bundle market. The denominator is the total quantity demanded for product h across all bundle markets.

Given predicted shares, we recover mean utilities via the BLP contraction mapping. Let s_{fht}^{obs} denote the observed market share, computed as quantity sold divided by total market quantity. The contraction mapping iterates as follows:

$$\delta_{fht}^{(r+1)} = \delta_{fht}^{(r)} + \log s_{fht}^{\text{obs}} - \log s_{fht}^{\text{pred}}(\boldsymbol{\theta}, \boldsymbol{\delta}^{(r)})$$

until convergence. We normalize the mean utility of one variety per product-month to zero.

The outer loop searches over $\boldsymbol{\theta}$ to maximize the likelihood of observed purchase decisions. Since the inner loop recovers mean utilities as a function of $\boldsymbol{\theta}$, the likelihood is ultimately a function of $\boldsymbol{\theta}$ alone. Each buyer i in bundle market (b, t) chooses a seller combination \mathbf{f}_{ibt}^* ; our demand model assigns probability $s_{i\mathbf{f}_{ibt}^*}$ to this choice. The

log-likelihood sums over all buyers and markets:

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{b,t} \sum_{i \in I_{bt}} \log s_{if_{bt}^*}(\boldsymbol{\theta}, \boldsymbol{\delta}(\boldsymbol{\theta}))$$

We maximize this likelihood to find the preference parameters that best explain observed purchase decisions.

5.1.2 Step 2: Price Coefficient

The mean utilities from Step 1 decompose into price and unobserved quality: $\delta_{fht} = \alpha \cdot p_{fht} + \xi_{fht}$.⁷ In our baseline specification, we calibrate $\alpha = -10$.

5.1.3 Step 3: Marginal Costs

To recover marginal costs, we invert the first-order conditions of the pricing game (Equation 9). For each firm f in period t :

$$\mathbf{mc}_{ft} = \mathbf{p}_{ft} - \boldsymbol{\Omega}_{ft}^{-1} \mathbf{Q}_{ft} / \alpha$$

Prices and quantities are observed in the data; $\boldsymbol{\Omega}_{ft}$ depends on estimated parameters from Step 1, and α is the price coefficient from Step 2. This allows us to express marginal costs as a function of observables and estimated parameters.⁸

⁷We normalize prices within each HSN category by dividing by the median price across all seller-month observations. Correspondingly, quantity is recomputed as transaction value divided by normalized price, yielding a value-equivalent volume measure. This normalization ensures cross-category comparability.

⁸For products with small market shares, the inversion can produce implausible marginal costs—negative values or extreme markups. We address this by imputing marginal costs for such observations. We first compute the Lerner index $L_{fht} = (p_{fht} - mc_{fht})/p_{fht}$ for all observations. We then regress the logit-transformed Lerner index on two predictors: the within-HSN decile of firm f 's sales value, and an indicator for whether h is the firm's main product. For observations with $mc \leq 0$ or $L \geq 0.95$, we impute marginal costs from the regression-predicted Lerner indices.

Table 5: Estimation Results

(a) Estimated Parameters		(b) Marginal Costs	
Parameter	Estimate	Statistic	Value
Demand Parameters		Mean MC	0.786
Distance (log)	-0.689	Median MC	0.871
Repeat Seller	9.123	% Imputed	15.529
Single-Seller Bonus	1.196	N Observations	35,945
Fixed Cost Parameters			
μ	40.425		
σ	7.090		

Notes. Left panel reports estimated demand and fixed cost parameters. Right panel reports summary statistics for recovered marginal costs.

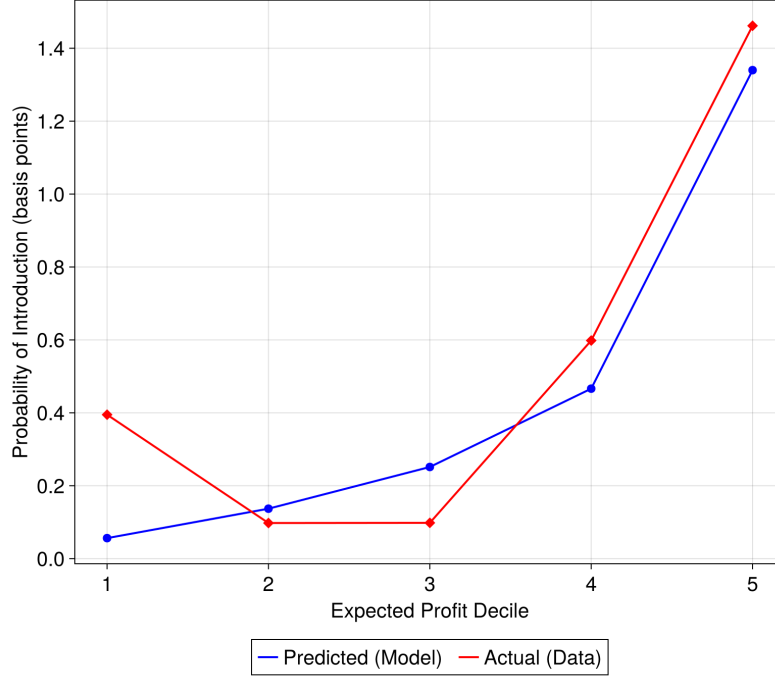
5.2 Fixed Costs

Recall that a firm introduces product h if expected gain in profits exceeds the fixed cost draw. Computing expected profit gain requires addressing two sources of uncertainty: the entrant does not know what quality and cost it will realize, and entry changes equilibrium prices.

We address both through simulation. For each potential entrant, we take 500 draws of (ξ_{fh}, mc_{fh}) from the empirical distribution of incumbents selling product h . For each draw, we add the entrant to the market, solve for post-entry Bertrand-Nash equilibrium prices, and compute the change in profit relative to the pre-entry equilibrium. Expected profit gain is the average across draws.⁹

⁹Computing expected profits is costly because each potential entrant requires solving for equilibrium prices

Figure 5: *Model Fit: Predicted vs. Actual Introduction Probabilities*



Notes. Predicted (model) and actual (data) product introduction probabilities by expected profit decile. Predicted probabilities computed using estimated fixed cost distribution.

Given expected change in profit $E[\Delta\Pi_f(h)]$, we estimate the fixed cost distribution by maximum likelihood. Since fixed costs are lognormal, the probability of entry is:

$$\Pr(\text{entry}_{fh}) = \Phi\left(\frac{\log E[\Delta\Pi_f(h)] - \mu}{\sigma}\right)$$

where Φ is the standard normal CDF and (μ, σ) parameterize the lognormal distribution.

500 times. To reduce computational burden, we employ choice-based sampling: since entry is rare, we include all observed entrants and a random sample of non-entrants, weighting by inverse sampling probabilities to correct for oversampling.

5.3 Results

Table 5a reports the estimated demand parameters. Distance enters negatively, while the repeat-seller and single-seller bonuses are positive: buyers prefer nearby, familiar sellers¹⁰ and value bundling from a single source. Table 5b summarizes the recovered marginal costs. With prices normalized so the median equals one, the median marginal cost is 0.87. We impute marginal costs for 15.5 percent of observations following the procedure described above in footnote 8.

To assess model fit, Figure 5 plots product introduction probabilities by quintile of expected profit gain. Expected profit gain is the model-implied change in firm profits from introducing a product: post-entry profit minus pre-entry profit. The figure reveals two patterns that support our model estimates. First, actual introduction probabilities increase with expected profit gain. Firms enter where the demand model predicts higher returns, and since expected profits depend on the estimated preference parameters, this upward slope confirms that distance, relationships, and bundling capture economically meaningful variation in profitability. Second, predicted probabilities closely track actual probabilities across quintiles, suggesting that the estimated fixed cost distribution rationalizes observed entry rates.

6 Counterfactual Analysis

The descriptive evidence in Section 3 showed that firms rely heavily on existing clients when selling new products, and that product introductions increase sales of complementary products already in the firm’s portfolio. The structural model formalized

¹⁰Appendix D reports an extended specification that allows the repeat-seller effect to differ for first-time bundle purchases. This helps assess whether the baseline repeat-seller coefficient reflects a relationship-based advantage or only buyer-seller-good-specific latent match quality. We find that the repeat-seller premium remains positive, and is larger, for first-time bundle purchases, which strengthens the relationship-based interpretation of β_r .

these patterns through two parameters: the repeat-seller bonus β_r , which captures the advantage of established buyer relationships, and the single-seller bonus γ , which captures the value of sourcing multiple products from a single supplier. We now use the estimated model to quantify how much these demand-side mechanisms matter for product introduction decisions.

We consider two counterfactual experiments. The first reduces β_r by 25 percent, simulating an environment where existing client relationships confer smaller advantages. The second reduces γ by 50 percent, simulating an environment where single-sourcing confers smaller advantages. For each experiment, we re-solve Bertrand-Nash equilibrium prices, recompute expected profits from entry, and compare the resulting introduction probabilities to the baseline.


6.1 Reducing Repeat-Seller Bonus

The parameter β_r captures several forces that make buyers prefer purchasing from familiar sellers: the cost of searching for new suppliers, the difficulty of evaluating unfamiliar sellers, and the risk of contracting with unknown parties. We do not distinguish among these channels. In our first counterfactual, we reduce β_r by 25 percent, simulating an environment where some combination of these forces is weaker, for example through online platforms that reduce search costs, quality certification that eases evaluation, or repeated interactions that build trust more quickly.

Reducing β_r raises product introduction probabilities. The left panel of Figure 6 shows that mean introduction probability rises from 0.50 to 0.74 basis points, an increase of 0.24 basis points.¹¹ The right panel plots median introduction probability

¹¹For each potential introduction (firm-product-month), we compute expected profit by simulating entry. The firm draws unobserved characteristics for its new product (quality and marginal cost) from the distribution of incumbents and we solve for post-entry Bertrand-Nash equilibrium prices. We average over 20 such draws to obtain expected profit.

Figure 6: *Impact on Product Entry from a 25% Reduction in Repeat-Seller Bonus*



Notes. The left panel compares aggregate introduction probability under baseline and counterfactual parameters. The right panel plots median introduction probability by decile of baseline expected profit. Weights account for choice-based sampling.

by decile of baseline expected profit. The absolute gap between the baseline and counterfactual curves widens from near zero at Q1 to approximately 0.45 basis points at Q10. This pattern reflects the convexity of the entry probability function. Because fixed costs are lognormally distributed, entry probability is a convex function of expected profit in the region where entry rates are low. Firms in higher profit deciles sit on a steeper portion of this curve, so a given improvement in expected profitability translates into a larger increase in entry probability.

Why does reducing β_r encourage entry? The repeat-seller bonus gives firms a competitive advantage with buyers they have previously transacted with. Incumbents have accumulated buyer relationships over time, which insulates them from competition. Potential entrants may have transacted with some of the same buyers in other product markets, but they have fewer buyer relationships in the market they are entering. Reducing β_r weakens this asymmetric advantage, so more entrants find entry profitable.

We explore this asymmetry directly. For each firm-product-month, we compute


revenue dependence, defined as the share of expected revenue derived from repeat buyers.¹² This measure captures how much a firm’s expected revenue depends on established buyer relationships, and therefore how much of its competitive position rests on the repeat-seller advantage. The argument above predicts that reducing β_r should favor firms with low revenue dependence. Since entrants have had fewer opportunities to accumulate buyer relationships, they should fall disproportionately into this group.

Figure 7 decomposes the effect of reducing β_r on incumbent outcomes. Mean price and markup responses are negative throughout the revenue-dependence distribution, with larger declines for firms that rely more heavily on repeat buyers. A lower β_r weakens relationship-based differentiation, making buyers more willing to substitute across sellers and thereby intensifying price competition. This effect is strongest for firms whose baseline sales depend most on repeat buyers, since they lose more of the relationship-based differentiation they previously enjoyed. The response of sales quantities changes sign across the revenue-dependence distribution. Because the model features only business stealing, lowering β_r reallocates a fixed pool of demand across sellers rather than expanding total demand. Firms with low revenue dependence gain sales on average as buyers become less attached to rival sellers, while firms with high revenue dependence lose sales on average as demand from their own repeat buyers erodes, though equilibrium price cuts partly offset this reallocation.

The price and quantity responses jointly determine the impact on profits. The left panel of Figure 8 plots mean log profit changes against revenue dependence. Incumbents with low revenue dependence gain: although their prices and markups fall, the business-stealing gains in quantity are large enough to raise profits (mean

¹²We compute revenue dependence separately for each firm-product-month. A buyer is classified as a repeat buyer of firm f in month t if their firm purchased any product from f prior to month t . For each buyer in the market, we use the estimated demand model to predict their probability of choosing firm f , weight by expected quantity, and sum. Revenue dependence is the ratio of this sum for repeat buyers to the sum across all buyers.


Figure 7: *Incumbent Quantity, Price, and Markup Effects from a 25% Reduction in Repeat-Seller Bonus*



Notes. Each panel plots the mean change in the outcome variable (log) against revenue dependence, defined as the share of a firm’s expected revenue from repeat buyers. Each point is a trimmed mean (1% tails excluded) within bins of revenue dependence. The markup panel restricts to observations with positive baseline and counterfactual markups.

log profit change of +0.12 in the lowest bin). For incumbents with high revenue dependence, the opposite holds. They lose profits on average because lower markups are compounded by losses in sales quantity as their repeat-seller advantage erodes. The mean profit response crosses zero at approximately 0.78 revenue dependence.

Figure 8: *Incumbent Profit Effects and Revenue Dependence Distribution*



Notes. The left panel plots the mean change in incumbent profit (log) against revenue dependence. Each point is a trimmed mean (1% tails excluded) within bins of revenue dependence. The right panel compares the distribution of revenue dependence for incumbents and potential entrants in the same product-month markets, computed under baseline parameters.

The right panel of Figure 8 shows why this matters for entry. It compares the distributions of revenue dependence for incumbents and potential entrants, computed under baseline parameters for the same product-month combinations. Two patterns emerge. First, the CDF for entrants lies to the left of the CDF for incumbents, confirming that entrants have systematically lower revenue dependence.¹³ Second, most incumbents fall above the 0.78 threshold. Reducing β_r therefore systematically lowers profits for incumbents while raising expected profits for potential entrants, and this asymmetry drives the increase in entry probabilities shown in Figure 6.

¹³Revenue dependence is not zero for potential entrants because firms entering a new product may already sell other products and therefore have buyer relationships with some of the same buyers.

6.2 Reducing Single-Seller Bonus

The single-seller bonus γ captures the value buyers place on sourcing multiple products from a single supplier. This preference may reflect reduced transaction costs, simplified logistics, or the overhead of managing multiple vendor relationships. We do not separately identify these channels. In our second counterfactual, we reduce γ by 50 percent, simulating an environment where single-sourcing confers smaller benefits, for example through logistics platforms that make multi-supplier coordination easier or procurement systems that reduce the cost of managing multiple vendors.

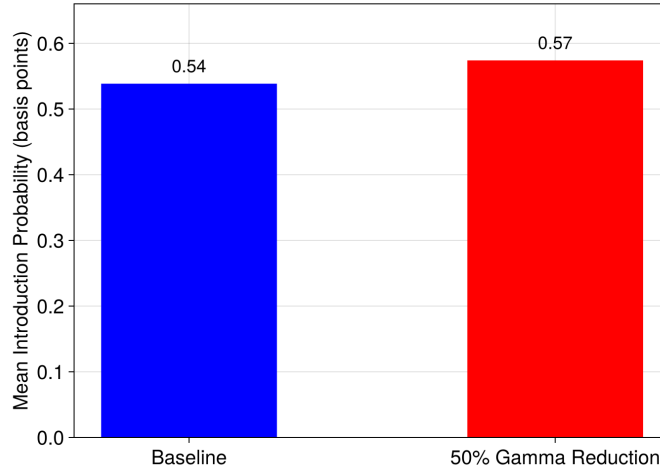
When γ falls, purchase decisions depend more on the individual merits of each product and less on whether a single firm can supply the entire bundle. Figure 9 shows that mean introduction probability rises from 0.54¹⁴ to 0.57 basis points, an increase of 0.03 basis points. The aggregate effect is smaller than in the repeat-seller counterfactual for two reasons. First, γ affects only multi-product purchases: buyers choosing whether to source a bundle from one firm or several. Single-product purchases, which account for a large share of transactions, are entirely unaffected. The repeat-seller bonus β_r , by contrast, enters every purchase. Second, the estimated γ of 1.20 is an order of magnitude smaller than β_r of 9.12 in utility terms. Even with a proportionally larger reduction (50 versus 25 percent), the resulting change in utility is smaller.

While the aggregate effect may be modest, it masks substantial heterogeneity across markets. To explore this, we define *bundling dependence* as the share of a firm's expected revenue from multi-product bundles where the firm supplies all products.¹⁵

¹⁴Baseline entry probabilities differ slightly across the two counterfactuals (0.50 vs. 0.54 basis points) because each experiment uses independent simulation draws to compute expected profits.

¹⁵We compute bundling dependence separately for each firm-product-month. For each buyer in a multi-product bundle market ($K \geq 2$), we use the estimated demand model to predict the probability of each seller combination and the associated revenue. Bundling dependence is the ratio of revenue from combinations where the firm supplies all products to total revenue from multi-product bundles.

Figure 9: *Impact on Product Entry from a 50% Reduction in Single-Seller Bonus*



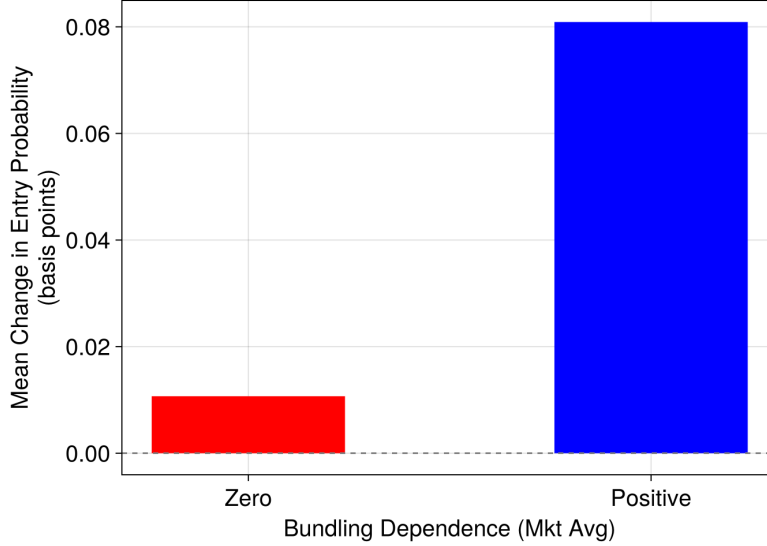
Notes. Aggregate introduction probability under baseline parameters and after reducing the single-seller bonus by 50 percent. Weights account for choice-based sampling.

This measure captures how much of a firm’s competitive position in bundle markets depends on its ability to supply complete bundles. We aggregate to the market level by computing the revenue-weighted average of bundling dependence among incumbents in each product-month market.

Figure 10 splits markets into those with zero and positive average incumbent bundling dependence. Entry probability gains are roughly eight times larger in positive-bundling-dependence markets (approximately 0.08 basis points) than in zero-bundling-dependence markets (approximately 0.01 basis points). Reducing γ erodes incumbent advantages precisely where those advantages exist. The small but nonzero gain in zero-bundling-dependence markets reflects equilibrium spillovers. Because the Bertrand-Nash equilibrium is solved jointly, price adjustments by firms that compete in both zero-bundling and positive-bundling markets can feed back into the focal market.

Figure 11 decomposes the effects on incumbent quantities, prices, markups, and

Figure 10: *Entry Probability Gains by Market-Level Bundling Dependence*



Notes. Entry probability gains by market-level bundling dependence, defined as the revenue-weighted average bundling dependence among incumbents in each product-month market. Bundling dependence measures the share of a firm's expected revenue from multi-product bundles where the firm supplies all products.

profits. The quantity response is monotonically decreasing in bundling dependence: firms with low bundling dependence gain sales (up to +0.05 log points at bundling dependence 0.1), while firms with high bundling dependence lose sales (-0.09 log points at 0.95), with the response crossing zero at approximately 0.45. As the single-seller bonus weakens, buyers substitute away from bundled purchases toward mixed-seller combinations, redistributing demand from high-bundling firms to those whose sales rest less on the ability to supply complete bundles.


The price response is nearly zero throughout the distribution, with all points within ± 0.01 log points. This near-absence of price adjustment distinguishes the γ channel from the β_r channel, where prices fall by 0.02 to 0.08 log points (Figure 7). Because γ enters utility only for multi-product bundles sourced from one firm, reducing it does not erode differentiation as broadly as reducing β_r . The γ channel

operates primarily through quantity reallocation rather than price competition.

The markup response is hump-shaped, peaking at approximately $+0.045$ log points at bundling dependence 0.5 and falling to -0.06 log points at 0.95, with the zero-crossing between bundling dependence 0.85 and 0.95. The shape reflects how the reduction in γ redistributes competitive advantage unevenly across firms. Firms with moderate bundling dependence benefit because their rivals at the top of the distribution lose more from the policy change than they do. These high-bundling-dependence rivals previously captured demand through the single-seller bonus; weakening that bonus erodes their competitive position, softens the competitive pressure they exert, and allows moderately bundling-dependent firms to charge slightly higher markups. For firms at the top of the bundling dependence distribution, the effect reverses. These firms lose their own bundling advantage directly. Buyers can more easily split purchases across sellers, making these firms' demand more price-sensitive, and their markups fall. Firms with low bundling dependence barely participate in bundle markets, so the reduction has little direct effect on their markups.

The profit response combines the quantity and markup channels. At low levels, large quantity gains drive profits upward ($+0.06$ log points at 0.1, $+0.065$ at 0.3), while the markup effect remains close to zero and contributes little in either direction. At moderate bundling dependence, markups peak but quantity gains have already eroded, so the two channels roughly offset. Beyond approximately 0.65 bundling dependence, change in profit turns negative: quantity losses are now large, and markups, rather than providing a counterweight, reinforce them. At 0.95, the two channels compound to produce a profit decline of -0.10 log points. The concentration of incumbent profit losses at high bundling dependence is consistent with the heterogeneity in entry gains shown in Figure 10: entry rises most in exactly the markets where incumbents lose the most.

Figure 11: *Incumbent Quantity, Price, Markup, and Profit Effects from a 50% Reduction in Single-Seller Bonus*



Notes. Each panel plots the mean change in the outcome variable (log) against bundling dependence, defined as the share of a firm's expected revenue from multi-product bundles where the firm supplies all products. Each point is a trimmed mean (1% tails excluded) within bins of bundling dependence. The bottom-left panel (markups) restricts to observations with positive baseline and counterfactual markups.

6.3 Discussion

The above two counterfactuals quantify two demand-side mechanisms that shape product scope. Buyers prefer purchasing from sellers they have previously transacted

with and value the convenience of sourcing multiple products from a single supplier. These preferences favor firms with established buyer relationships and broad product portfolios, making it difficult for firms without these assets to introduce new products. Reducing repeat-seller preferences or single-sourcing preferences weakens the position of established incumbents and increases entry rates for potential new products.

To gauge the economic significance of these entry probability changes, we translate them into approximate counts of additional product introductions. Our estimation sample contains roughly 501,000 weighted potential entrant opportunities (firm-product-month observations) over 11 months, or approximately 547,000 annualized.¹⁶ The 25 percent reduction in β_r raises mean entry probability by 0.24 basis points, implying roughly 13 additional product introductions per year. For context, the data contain approximately 28 product introductions per year, so this counterfactual represents a roughly 50 percent increase over the observed rate. The 50 percent reduction in γ raises mean entry probability by 0.03 basis points, implying roughly 2 additional product introductions per year, an approximately 8 percent increase over the observed rate.

These two counterfactuals use different percentage reductions (25 versus 50 percent) and operate through different channels, so the magnitudes are not directly comparable. The experiments also rule out dynamic responses: incumbents do not adjust portfolios in response to new entry, and entrants do not condition on the sequential arrival of competitors. Our estimates therefore isolate the static, partial-equilibrium effects of each parameter change on the entry margin. Although these are static, partial-equilibrium exercises, the counterfactuals map to concrete policy interventions. Reducing the repeat-seller advantage corresponds to platforms or institutions that lower search costs and ease supplier evaluation: verified supplier ratings, quality

¹⁶The estimation sample contains 1,026 potential entrant observations (26 actual entrants and 1,000 non-entrants drawn with choice-based sampling). Each non-entrant receives a weight of approximately 501 to reflect the size of the non-entrant population. The approximate number of additional introductions equals the change in mean entry probability multiplied by total weighted opportunities.

certification programs, and digital marketplaces that help buyers discover and assess unfamiliar sellers. Reducing the single-seller bonus corresponds to logistics infrastructure and procurement systems that lower the cost of coordinating across multiple suppliers. Both types of intervention are active areas of platform design in Indian B2B e-commerce.¹⁷

7 Conclusion

This paper studies how demand-side factors shape firms' decisions to expand their product portfolios. Using administrative transaction records from India's e-way bill system, we document three empirical patterns: firms rely on existing clients for over half of new product sales; introducing a product increases sales of complementary products already in the firm's portfolio; and market access to existing clients, particularly for complements, predicts which products firms introduce. We develop a structural model that formalizes these patterns through two mechanisms: a repeat-seller preference that captures the advantage of established relationships, and a single-seller bonus that captures buyers' preference for sourcing multiple products from the same firm. The estimated model fits observed introduction probabilities well and allows us to quantify the role of each mechanism.

Counterfactual exercises reveal that both mechanisms affect product introduction, though to different degrees. Reducing the repeat-seller preference by 25 percent increases mean introduction probability by 50 percent. This is because incumbents have accumulated buyer relationships over time; potential entrants have not. Thus, when the repeat-seller preference falls, this incumbency advantage weakens, and product

¹⁷For example, IndiaMART's TrustSEAL program provides third-party verification of supplier credentials to reduce buyer search and evaluation costs (<https://corporate.indiamart.com/mdc-ts-combo/>), while Udaan's integrated logistics network consolidates multi-supplier fulfillment (<https://udaan.com/about-us>) and Moglix's procurement platform manages sourcing and delivery across thousands of industrial suppliers (<https://business.moglix.com/>).

introduction probability rises. Reducing the single-seller bundling bonus by 50 percent produces a smaller aggregate effect, though gains concentrate in markets where incumbents rely heavily on bundling. These results show that buyer preferences over supplier relationships and multi-product sourcing create barriers to product introduction that favor established firms with broad portfolios and deep customer bases.

These findings also speak to firm growth in developing economies. They suggest that the constraint on product variety may not just be a lack of productive capacity or entrepreneurial talent, but rather the difficulty of breaking into established buyer-seller networks. In settings where formal contracting is costly and information about suppliers is limited, buyer-seller relationships play a large role in shaping which firms can expand. Our results show that firms without established customer networks face barriers when introducing new products. Understanding these demand-side forces complements supply-side explanations and helps explain why product scope remains concentrated among a small number of firms.

References

- Atkin, D., Khandelwal, A. K., & Osman, A. (2017). Exporting and firm performance: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, *132*(2), 551–615.
- Banerjee, A. V., & Duflo, E. (2014). Do firms want to borrow more? Testing credit constraints using a directed lending program. *The Review of Economic Studies*, *81*(2), 572–607.
- Bernard, A. B., Blanchard, E. J., Van Beveren, I., & Vandebussche, H. (2019). Carry-along trade. *The Review of Economic Studies*, *86*(2), 526–563.
- Bernard, A. B., Redding, S. J., & Schott, P. K. (2010). Multiple-product firms and product switching. *American Economic Review*, *100*(1), 70–97.
- Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica : journal of the Econometric Society*, *63*(4), 841.
- Besley, T., & Burgess, R. (2004). Can labor regulation hinder economic performance? Evidence from India. *The Quarterly Journal of Economics*, *119*(1), 91–134.
- Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., & Roberts, J. (2013). Does management matter? Evidence from India. *The Quarterly Journal of Economics*, *128*(1), 1–51.
- Boehm, J., Dhingra, S., & Morrow, J. (2022). The comparative advantage of firms. *Journal of Political Economy*, *130*(12), 3025–3100.
- Buera, F. J., Kaboski, J. P., & Shin, Y. (2011). Finance and development: A tale of two sectors. *American Economic Review*, *101*(5), 1964–2002.
- Callaway, B., & Sant’Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.
- Eaton, J., Eslava, M., Jinkins, D., Krizan, C. J., & Tybout, J. R. (2025). A search and learning model of export dynamics. *Journal of International Economics*, *157*(100).
- Eckel, C., & Neary, J. P. (2010). Multi-product firms and flexible manufacturing in the global economy. *Review of Economic Studies*, *77*(1), 188–217.

- Foster, L., Haltiwanger, J., & Syverson, C. (2016). The slow growth of new plants: Learning about demand? *Economica*, 83(329), 91–129.
- Gentzkow, M. (2007). Valuing New Goods in a Model with Complementarity: Online Newspapers. *The American Economic Review*, 97(3), 713–744.
- Goldberg, P. K., Khandelwal, A. K., Pavcnik, N., & Topalova, P. (2010). Imported intermediate inputs and domestic product growth: Evidence from India. *The Quarterly Journal of Economics*, 125(4), 1727–1767.
- Goldberg, P. K., & Reed, T. (2023). Presidential address: Demand-side constraints in development: The role of market size, trade, and (in)equality. *Econometrica*, 91(6), 1915–1950.
- Gourio, F., & Rudanko, L. (2014). Customer capital. *The Review of Economic Studies*, 81(3), 1102–1136.
- Hardy, M., & Kagy, G. (2020). It’s getting crowded in here: Experimental evidence of demand constraints in the gender profit gap. *The Economic Journal*, 130(631), 2272–2290.
- Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and manufacturing TFP in China and India. *The Quarterly Journal of Economics*, 124(4), 1403–1448.
- Leone, F., Macchiavello, R., & Reed, T. (2025). The high and falling price of cement in africa. *American Economic Journal: Applied Economics*, 17(2), 1–40.
- Macchiavello, R. (2022). Relational contracts and development. *Annual Review of Economics*, 14, 337–362.
- Macchiavello, R., & Morjaria, A. (2015). The value of relationships: Evidence from a supply shock to Kenyan rose exports. *American Economic Review*, 105(9), 2911–2945.
- McMillan, J., & Woodruff, C. (1999). Interfirm relationships and informal credit in Vietnam. *The Quarterly Journal of Economics*, 114(4), 1285–1320.
- Monarch, R. (2022). It’s not you, it’s me: Prices, quality, and switching in u.s.-china trade relationships. *The Review of Economics and Statistics*, 104(5), 909–928.

Press Information Bureau. (2021, April 1). *7.12 crores e-way bills generated in march 2021, highest ever since launch of e-way bill system*. Government of India, Ministry of Finance.

Startz, M. (2024). *The value of face-to-face: Search and contracting problems in Nigerian trade*.

APPENDIX

Customer-Driven Product Scope

Contents

A Data Construction Details	3
A.1 E-Way Bill System Details	3
A.2 Raw Data Characteristics	3
A.3 HSN Product Classification	4
A.4 Sample Construction	4
A.5 Reduced-Form Analysis Sample	5
B Stylized Model	8
B.1 Setup	8
B.2 Choice Regions	9
B.3 Predictions	9
C Derivation of First-Order Conditions	12
C.1 Choice Probability Derivatives	12

C.2	First-Order Condition	13
C.3	Deriving the Ω -Matrix System	14
D	Extended Demand Specification	16

A Data Construction Details

This appendix describes the e-way bill system, the raw data, and our sample construction.

A.1 E-Way Bill System Details

Each e-way bill consists of two parts. Part A records transaction details: the GSTIN and PIN code of buyer and seller, the invoice number, date, and value, and the HSN code identifying the product. Part B captures transportation details: the vehicle number, mode of transport, and distance traveled. A separate e-way bill is required for each invoice, so a single truck carrying goods covered by multiple invoices generates multiple e-way bills.

The system includes enforcement mechanisms at highway checkpoints. GST officers can intercept any vehicle and verify that its cargo matches the e-way bill documentation. Failure to produce a valid e-way bill results in penalties of Rs. 10,000 or the tax evaded, whichever is higher, and the goods and vehicle may be detained. In the first three years alone (April 2018 to March 2021), over 1.8 billion e-way bills were generated (Press Information Bureau, [2021](#)). This reflects widespread adoption among formal sector firms.

A.2 Raw Data Characteristics

The raw data span 24 monthly files containing approximately 439 million transaction records for the period April 2018 through March 2020. Each record corresponds to a single e-way bill and includes:

- GSTIN and PIN code of consignor (seller)
- GSTIN and PIN code of consignee (buyer)
- Invoice number, date, and value
- HSN code and product description
- Quantity and unit of measurement
- Transport mode and vehicle number
- Distance traveled (kilometers)

A.3 HSN Product Classification

The Harmonized System of Nomenclature (HSN) is an international standard developed by the World Customs Organization. HSN codes identify traded products at varying levels of granularity: the first two digits indicate the chapter (broad category), the next two the heading, and subsequent digits provide finer detail.

Indian GST requires businesses with annual turnover above Rs. 5 crore (600K USD) to report at least four HSN digits. The most common categories include cement (HSN 2523), pharmaceuticals (HSN 3004), auto parts (HSN 8708), and ceramic tiles (HSN 6907).

A.4 Sample Construction

We apply the following filters to the raw data:

1. **Restrict time period:** Exclude transactions dated on or after April 1, 2020, to avoid COVID-19 disruptions.

2. **Remove intra-firm transactions:** Remove shipments between establishments of the same firm (where buyer and seller share a tax identifier).
3. **Retain valid HSN codes:** Drop records with missing or invalid HSN codes. Codes shorter than four digits after standardization account for 5.4% of transactions but only 0.4% of total value.
4. **Clean sales and quantities:** Exclude records with zero or negative quantities or values.
5. Winsorize transaction values at the 99.999th percentile.

After these filters, the cleaned data contain 53.8 million transactions involving 242,474 unique selling firms and 649,503 unique buying firms.

Additional restrictions for main analysis. We impose two further restrictions to focus on established firms with regular activity:

1. **Firm activity:** Retain only firms active for at least 13 months over the sample period. This filter reduces the seller count from 242,474 to 39,964 (16.5% retained).
2. **Product activity:** Retain only seller-product pairs with more than 50 transactions over the sample period. This filter removes sporadic or occasional product lines.

Table [A.1](#) summarizes sample attrition at each step.

A.5 Reduced-Form Analysis Sample

For the reduced-form event study analysis in Section [3](#), we aggregate products to the 4-digit HSN level.

Table A.1: *Sample Attrition*

Step	Transactions	Sellers	% Retained
Raw data (Apr 2018–Mar 2020)	439,000,000	—	—
After all cleaning filters	53,800,000	242,474	12.3%
After firm activity filter (>12 months)	46,100,000	39,964	85.7%
After product activity filter (>50 txns)	42,600,000	39,344	92.4%

Notes. This table shows the number of transactions and unique sellers at each stage of sample construction. The cleaning filters include: date restriction (pre-COVID), intra-firm removal, HSN validation, value/quantity cleaning, and outlier winsorization. Percentages in the final column show retention relative to the previous step.

Table A.2 presents summary statistics for this sample. The resulting panel contains 1.70 million seller-product-month observations across 39,964 firms and 1,029 product categories. The median firm-product-month observation records Rs. 4.60 lakh in sales across 7 transactions to 4 distinct buyers. The mean is substantially higher (Rs. 49 lakh per firm-product-month), reflecting a right-skewed distribution driven by differences in firm size.

Under the product introduction definition described in Section 2, this sample yields 2,368 product introductions across 1,300 firms, representing 2.7% of all firm-product pairs. The typical introduction occurs 9 months after the firm’s entry into the data. Most firms that introduce products add just one or two new categories (median 1, mean 1.8), though a few expand into dozens.

Table A.2: *Sample Summary Statistics (4-Digit HSN)*

Variable	Value
Total Observations	42,602,732
Unique Seller Firms	39,344
Unique Buyer Firms	120,633
Unique Products (HSN4)	1,116
Unique E-way Bills	24,963,653
Total Transaction Value (Billion USD)	97.670
Median Transaction Value (Hundred USD)	8.160
Median Transactions per Seller	175
Median Products per Seller	1

Notes. This table presents descriptive statistics for the reduced-form analysis sample, aggregated to the 4-digit HSN level. The sample includes only firms that were active for at least 13 months and excludes within-firm shipments. Transaction values are in Indian Rupees (1 USD \approx 85 INR).

B Stylized Model

This appendix develops a stylized model of product introduction that generates testable predictions for our reduced-form analysis. The model incorporates two key features of demand. First, existing client relationships provide additional utility from repeat purchases. Second, buyers receive additional utility from purchasing products together, which leads to complementarity between products. We use this framework to derive predictions about what drives product introduction and what happens after introduction.

B.1 Setup

Consider a market with two products, A and B . A firm f currently sells product A and is deciding whether to introduce product B . Buyers, indexed by i , can purchase products individually or as a bundle. After introduction, buyers choose among four options: purchase only A , purchase only B , purchase both as a bundle, or purchase neither. Each buyer demands q_i units conditional on purchase. Prices are exogenous.

Buyer i 's utility from purchasing product A , product B , or the bundle AB is:

$$\begin{aligned}U_A^i &= \mu_f + \mu_A + \beta_r \cdot \mathbf{1}[\text{repeat}_{i,f}] + \beta_d \cdot \log d_{if} - \alpha \cdot p_A + \nu_A^i \\U_B^i &= \mu_f + \mu_B + \beta_r \cdot \mathbf{1}[\text{repeat}_{i,f}] + \beta_d \cdot \log d_{if} - \alpha \cdot p_B + \nu_B^i \\U_{AB}^i &= U_A^i + U_B^i + \Gamma \\U_0^i &= 0\end{aligned}$$

Utility from purchasing product A or B depends on firm and product characteristics μ_f , μ_A , μ_B , whether the buyer has an existing relationship with the firm, distance, and prices. The terms ν_A^i and ν_B^i capture idiosyncratic buyer preferences for each product.

The third equation defines utility when both products are purchased together as a bundle. Buyers receive utility from each product, $U_A^i + U_B^i$, plus an additional utility Γ . When $\Gamma > 0$, products are complements: buyers value the bundle more than the sum of its parts. This formulation follows Gentzkow (2007). The last equation normalizes the outside option to zero.

B.2 Choice Regions

Each buyer chooses the option with highest utility. Figure B.1 illustrates the choice regions in (U_A, U_B) space. When $\Gamma = 0$, buyer i purchases both products if and only if $U_A^i > 0$ and $U_B^i > 0$. When $\Gamma > 0$, the bundle region expands. A buyer may now purchase the bundle even when one product has negative standalone utility, provided the combined utility $U_A^i + U_B^i + \Gamma$ exceeds zero.

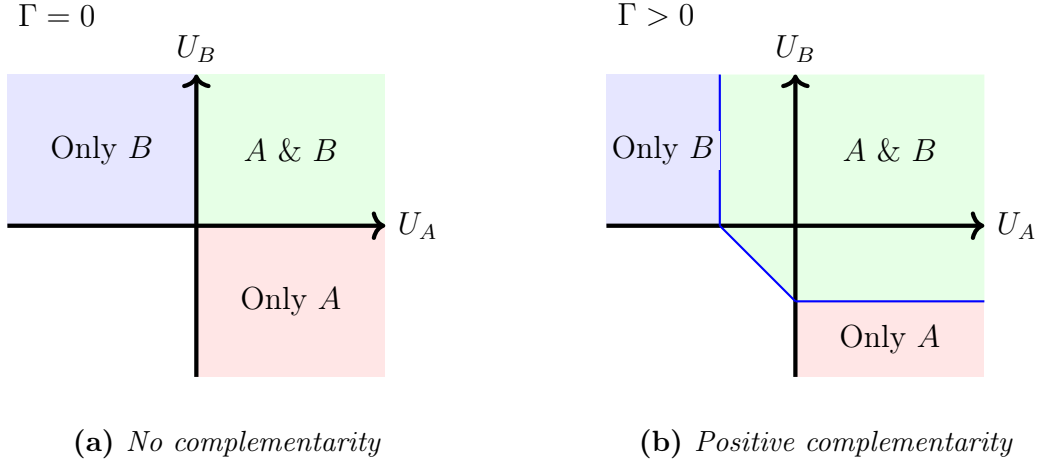
This expansion has implications for product introduction. Buyers with $U_A^i \in (-\Gamma, 0)$ would not purchase A alone, but may purchase it as part of the bundle. When $\Gamma > 0$, introducing B expands demand for A by converting these marginal non-buyers into bundle purchasers. This makes introducing B more profitable than when $\Gamma = 0$, and firms are more likely to introduce products that complement their existing offerings.

B.3 Predictions

We now derive testable predictions from this framework.

Prediction 1: Product introduction is driven by market access, with different roles for existing clients and new buyers. Consider first the case where $\Gamma = 0$. A firm is more likely to introduce a product when located near buyers with high demand for

Figure B.1: *Choice Regions Under Different Complementarity*



Notes. Choice regions in utility space. When $\Gamma = 0$, buyers purchase both products if and only if $U_A > 0$ and $U_B > 0$. When $\Gamma > 0$, the bundle region expands: buyers may purchase both products even when one has negative standalone utility, provided the combined utility exceeds the threshold.

that product. Expected sales of the new product B are:

$$Q_B = \int q_i \cdot \mathbf{1}\{U_B^i > 0\} dF(i)$$

With simplifying distributional assumptions, this expression yields:

$$\log(\text{Sales of } B) \approx \lambda_f + \lambda_B + \log \left(\theta^e \sum_{i \in \mathcal{E}} \frac{q_i}{d_{if}^{|\beta_d|}} + \theta^n \sum_{i \in \mathcal{N}} \frac{q_i}{d_{if}^{|\beta_d|}} \right)$$

where \mathcal{E} denotes existing clients and \mathcal{N} denotes new buyers. The summation terms capture distance-weighted demand potential, which we call *market access*. Firms with high-demand buyers nearby—high q_i and low d_{if} —have higher market access and are more likely to introduce. Existing clients and new buyers contribute differently to expected sales. If existing relationships provide additional utility, we expect $\theta^e > \theta^n$: a given level of demand from existing clients translates into higher sales than the same demand from new buyers.

Prediction 2: Product introduction increases sales of existing products, and com-

plementarity drives introduction. When $\Gamma > 0$, the increase in sales from introducing product B has three components:

$$\begin{aligned} \Delta Q = \int q_i \cdot & \left[\underbrace{\mathbf{1}\{U_B^i > 0\}}_{\text{sales of } B} \right. \\ & + \underbrace{\mathbf{1}\{0 > U_A^i > -\Gamma\} \cdot \mathbf{1}\{U_A^i + U_B^i > -\Gamma\}}_{\text{higher bundled sales of } A} \\ & \left. + \underbrace{\mathbf{1}\{0 > U_B^i > -\Gamma\} \cdot \mathbf{1}\{U_A^i + U_B^i > -\Gamma\}}_{\text{higher bundled sales of } B} \right] dF(i) \end{aligned}$$

The first term captures direct sales of B , as in Prediction 1. The second term captures buyers with $U_A^i \in (-\Gamma, 0)$ who would not purchase A alone but now purchase it as part of the bundle. The third term is symmetric for B . Introducing a new product therefore increases sales of existing products, and firms are more likely to introduce products that complement their existing offerings.

Prediction 3: The effect on prices of existing products is ambiguous. We explore this prediction through numerical simulation. We find a non-monotonic relationship between Γ and the price of existing products: prices can either increase or decrease after product introduction, depending on the degree of complementarity. This ambiguity arises from two opposing forces. First, lowering prices becomes more attractive. Before introduction, a price cut on A only increased sales of A . After introduction, it increases sales of both A and B via the bundle. Second, raising prices also becomes more attractive. Buyers who were previously marginal now receive additional utility from the bundle, making them inframarginal. The firm can raise prices without losing them, allowing it to extract surplus.

C Derivation of First-Order Conditions

This appendix derives the first-order conditions for pricing and expresses them in matrix form. The resulting system allows us to recover marginal costs from observed prices, quantities, and estimated demand parameters. Following the main text, we suppress time subscripts throughout. Recall that markets are defined by a bundle $b = \{h_1, \dots, h_K\}$, and buyers choose among seller combinations $\mathbf{f} = (f_1, \dots, f_K)$ with choice probabilities $s_{i\mathbf{f}} = \exp(V_{i\mathbf{f}}) / \sum_{\mathbf{f}'} \exp(V_{i\mathbf{f}'})$. Mean utility is $\delta_{fh} = \alpha \cdot p_{fh} + \xi_{fh}$, where $\alpha < 0$ is the price coefficient.

C.1 Choice Probability Derivatives

We first derive how prices affect choice probabilities. The key result expresses the derivative $\partial s_{i\mathbf{f}} / \partial p_{fh}$ in terms of the marginal probability S_{ifh} that buyer i purchases product h from firm f . This decomposition separates own-price effects from cross-price substitution.

Let $D = \sum_{\mathbf{f}'} \exp(V_{i\mathbf{f}'})$ denote the denominator, so that $s_{i\mathbf{f}} = \exp(V_{i\mathbf{f}}) / D$. The derivative is:

$$\begin{aligned} \frac{\partial s_{i\mathbf{f}}}{\partial p_{fh}} &= \frac{D \cdot \frac{\partial \exp(V_{i\mathbf{f}})}{\partial p_{fh}} - \exp(V_{i\mathbf{f}}) \cdot \frac{\partial D}{\partial p_{fh}}}{D^2} \\ &= \frac{\exp(V_{i\mathbf{f}})}{D} \cdot \frac{\partial V_{i\mathbf{f}}}{\partial p_{fh}} - \frac{\exp(V_{i\mathbf{f}})}{D} \sum_{\mathbf{f}'} \frac{\exp(V_{i\mathbf{f}'})}{D} \cdot \frac{\partial V_{i\mathbf{f}'}}{\partial p_{fh}} \\ &= s_{i\mathbf{f}} \left(\frac{\partial V_{i\mathbf{f}}}{\partial p_{fh}} - \sum_{\mathbf{f}'} s_{i\mathbf{f}'} \cdot \frac{\partial V_{i\mathbf{f}'}}{\partial p_{fh}} \right) \end{aligned}$$

From the utility specification, price p_{fh} only affects seller combinations where firm f supplies product h :

$$\frac{\partial V_{i\mathbf{f}}}{\partial p_{fh}} = \mathbf{1}[f_h = f] \cdot \alpha \cdot w_{ih}$$

where w_{ih} is the quantity weight. Substituting yields:

$$\begin{aligned}\frac{\partial s_{i\mathbf{f}}}{\partial p_{fh}} &= \alpha \cdot w_{ih} \cdot s_{i\mathbf{f}} \cdot \left(\mathbf{1}[f_h = f] - \sum_{\mathbf{f}'} s_{i\mathbf{f}'} \cdot \mathbf{1}[f'_h = f] \right) \\ &= \alpha \cdot w_{ih} \cdot s_{i\mathbf{f}} \cdot (\mathbf{1}[f_h = f] - S_{ifh})\end{aligned}$$

where $S_{ifh} = \sum_{\mathbf{f}'} s_{i\mathbf{f}'} \cdot \mathbf{1}[f'_h = f]$ is the marginal probability that buyer i purchases product h from firm f .

The term $(\mathbf{1}[f_h = f] - S_{ifh})$ determines the sign of the share response. Consider a price increase for variety (f, h) . Seller combinations that include firm f for product h become less attractive. For these combinations, the indicator equals one, so the term is positive. Since $\alpha < 0$, their shares decrease. The lost demand must go somewhere: seller combinations that do *not* include firm f gain share. For these combinations, the indicator is zero, so the term is negative. Since $\alpha < 0$, their shares increase.

C.2 First-Order Condition

We now derive the firm's optimality condition for pricing. Firm f 's profit and its derivative with respect to p_{fh} are:

$$\begin{aligned}\Pi_f &= \sum_{b \in \mathcal{B}} \sum_{i \in I_b} \sum_{\mathbf{f}' \in \mathcal{C}_b} s_{i\mathbf{f}'} \left(\sum_{k=1}^K \mathbf{1}[f'_k = f] \cdot (p_{fh_k} - mc_{fh_k}) \cdot q_{i,h_k} \right) \\ \frac{\partial \Pi_f}{\partial p_{fh}} &= \sum_{b \ni h} \sum_i \left[\underbrace{\sum_{\mathbf{f}'} s_{i\mathbf{f}'} \cdot \mathbf{1}[f'_h = f] \cdot q_{ih}}_{\text{Direct Revenue Effect}} + \underbrace{\sum_{\mathbf{f}'} \frac{\partial s_{i\mathbf{f}'}}{\partial p_{fh}} \left(\sum_{k \in b} \mathbf{1}[f'_k = f] (p_{fk} - mc_{fk}) q_{ik} \right)}_{\text{Demand Diversion Effect}} \right]\end{aligned}$$

The direct revenue effect captures the additional revenue from selling at a higher price, as each unit sold contributes more to profit. The demand diversion effect captures the cost of losing sales as buyers substitute away. Raising the price of product

h reduces demand for seller combinations that include the firm for h , lowering sales volume. For multi-product firms, there is an additional effect. Buyers who substitute away also stop purchasing the firm's other products in the same bundle. The inner sum over $k \in b$ captures the margin on each product the firm sells. This is multi-product pricing internalization, where the firm accounts for the impact of pricing h on its other products.

The direct revenue effect equals $Q_{fh} = \sum_{b \ni h} \sum_i S_{ifh} \cdot q_{ih}$, the total quantity of product h sold by firm f . Substituting the share derivative from above and setting $\partial \Pi_f / \partial p_{fh} = 0$ yields the first-order condition:

$$Q_{fh} + \sum_{b \ni h} \sum_i \alpha w_{ih} \sum_{f'} s_{if'} (\mathbf{1}[f'_h = f] - S_{ifh}) \left(\sum_{k \in b} \mathbf{1}[f'_k = f] (p_{fk} - mc_{fk}) q_{ik} \right) = 0$$

C.3 Deriving the Ω -Matrix System

The goal is to express the FOC in matrix form, where a single matrix Ω captures how each product's margin contributes to the optimality condition for every other product in the firm's portfolio. We proceed by reordering the summations to isolate margin terms.

Multiplying by $-1/\alpha$ and moving Q_{fh} to the right-hand side gives

$$\begin{aligned}
& \sum_{b \ni h} \sum_i w_{ih} \sum_{\mathbf{f}'} s_{i\mathbf{f}'} (S_{ifh} - \mathbf{1}[f'_h = f]) \left(\sum_{k \in b} \mathbf{1}[f'_k = f] (p_{fk} - mc_{fk}) q_{ik} \right) = \frac{Q_{fh}}{\alpha} \\
& \sum_{b \ni h} \sum_i \sum_{k \in b} (p_{fk} - mc_{fk}) \left(w_{ih} \cdot q_{ik} \sum_{\mathbf{f}'} s_{i\mathbf{f}'} (S_{ifh} - \mathbf{1}[f'_h = f]) \mathbf{1}[f'_k = f] \right) = \frac{Q_{fh}}{\alpha} \\
& \sum_{b \ni h} \sum_{k \in b} (p_{fk} - mc_{fk}) \left(\sum_i w_{ih} \cdot q_{ik} \sum_{\mathbf{f}'} s_{i\mathbf{f}'} (S_{ifh} - \mathbf{1}[f'_h = f]) \mathbf{1}[f'_k = f] \right) = \frac{Q_{fh}}{\alpha} \\
& \sum_{k \in \mathcal{H}_f} (p_{fk} - mc_{fk}) \underbrace{\left(\sum_{b \ni \{h,k\}} \sum_{i \in I_b} w_{ih} \cdot q_{ik} \sum_{\mathbf{f}'} s_{i\mathbf{f}'} (S_{ifh} - \mathbf{1}[f'_h = f]) \mathbf{1}[f'_k = f] \right)}_{\Omega_{hk}} = \frac{Q_{fh}}{\alpha}
\end{aligned}$$

The first line rearranges the FOC. The second swaps \mathbf{f}' and k to bring the margin outside. The third swaps i and k . The fourth swaps bundles and products, iterating over all k in \mathcal{H}_f and bundles containing both h and k . The underbrace defines the matrix element Ω_{hk} , and the FOC becomes

$$\sum_{k \in \mathcal{H}_f} \Omega_{hk} (p_{fk} - mc_{fk}) = Q_{fh}/\alpha$$

This holds for every product h in firm f 's portfolio. In matrix notation,

$$\mathbf{\Omega}(\mathbf{p} - \mathbf{mc}) = \mathbf{Q}/\alpha$$

which matches Equation 9 in the main text.

To illustrate, consider a firm selling products 1 and 2:

$$\Omega_{11}(p_1 - mc_1) + \Omega_{12}(p_2 - mc_2) = Q_1/\alpha$$

$$\Omega_{21}(p_1 - mc_1) + \Omega_{22}(p_2 - mc_2) = Q_2/\alpha$$

Consider the first equation. The term Ω_{11} captures the own-price effect—how raising p_1 affects demand for product 1. The term Ω_{12} captures the cross-price effect—how raising p_1 affects demand for product 2 through bundle substitution.

D Extended Demand Specification

A natural concern in our baseline demand specification is that the repeat-seller coefficient may capture persistent unobserved preference heterogeneity rather than state dependence. For example, some buyer-seller-good matches may be intrinsically good along dimensions we do not observe, in which case repeated purchases would reflect latent match quality rather than a causal effect of prior trade. If that were the case, the baseline repeat-seller coefficient would overstate the degree to which past transactions themselves make a seller more attractive.

To probe this concern, we focus on a setting where the distinction is especially sharp: first-time bundle purchases. Suppose the baseline repeat-seller effect were driven entirely by buyer-seller-good-specific latent match quality. Then prior purchases from the same seller in other goods should not create much advantage when the buyer is purchasing a bundle for the first time. By contrast, if prior transactions generate a seller-level relationship advantage through trust, learning, or lower search and coordination costs, then familiarity with the seller should remain valuable precisely in this unfamiliar purchase setting. This comparison cannot rule out all forms of unobserved heterogeneity, but it can rule out the narrower explanation that the repeat-seller effect is entirely good specific.

Motivated by this logic, we estimate an extended specification that stays close to the baseline model and adds only one extra interaction term:

$$\begin{aligned}
 V_{i\mathbf{f}} = & \sum_{k=1}^K w_{ik} \cdot \delta_{f_k, h_k} + \beta_r \sum_{k=1}^K w_{ik} \cdot \mathbf{1}[\text{repeat}_{i, f_k}] + \gamma \cdot \mathbf{1}[\text{single}] \\
 & + \beta_{rb} \left(\sum_{k=1}^K w_{ik} \cdot \mathbf{1}[\text{repeat}_{i, f_k}] \right) \mathbf{1}[\text{new bundle}_{ib}] \\
 & + \beta_d \cdot \log \left(\sum_{k=1}^K w_{ik} \cdot d_{i, f_k} \right).
 \end{aligned} \tag{11}$$

The indicator $\mathbf{1}[\text{new bundle}_{ib}]$ equals one if buyer i has not previously purchased bundle b before period t . Thus, β_{rb} captures whether the repeat-seller advantage is especially strong when the bundle itself is unfamiliar. In this specification, the repeat-seller effect for previously purchased bundles is β_r , while the repeat-seller effect for first-time bundle purchases is $\beta_r + \beta_{rb}$.

Table D.1 reports the estimates. Relative to the baseline specification (V1), the repeat-seller coefficient remains large in the extended specification, and the interaction term is positive. Since the repeat-seller effect equals β_r for previously purchased bundles and $\beta_r + \beta_{rb}$ for first-time bundle purchases, the positive interaction implies that seller familiarity is even more valuable when buyers are purchasing an unfamiliar bundle. This pattern is consistent with a relationship-based interpretation. When a buyer is procuring a new combination of goods, prior dealings with a seller may reduce uncertainty and lower search or coordination costs. By contrast, if persistence reflected only buyer-seller-good-specific latent match quality, prior purchases from the same seller in other goods would have limited power to explain why that seller is preferred for a first-time bundle purchase.

This exercise does not fully separate state dependence from all forms of persistent heterogeneity. In particular, it does not rule out buyer-seller match quality that carries across goods, such as time-invariant differences in reliability, delivery quality, credit terms, or informal contracting quality. If a buyer prefers a seller for these reasons across all goods, that seller-level match quality could still generate a positive repeat-seller premium even for first-time bundle purchases. At the same time, the extended specification sharpens the interpretation of the baseline results by showing that the repeat-seller premium is not confined to a single good-specific match, but remains strong even when buyers are purchasing a new bundle. We therefore view this exercise as strengthening the relationship-based interpretation of repeat purchases, while recognizing that the estimated repeat-seller term may still reflect a combination

Table D.1: *Demand Parameter Estimates: Baseline vs. Extended Specification*

	V1	V2
Distance (log)	-0.689	-0.696
Repeat Seller	9.123	8.985
Single-Seller Bonus	1.196	0.850
Repeat Seller \times New Bundle	—	3.265

Notes. V1 is the baseline specification with three preference parameters. V2 adds an interaction between the weighted repeat-seller term and an indicator for whether the buyer is purchasing the bundle for the first time. Both specifications are estimated by maximum likelihood using the nested fixed-point procedure described in Section 5.

of true state dependence and persistent buyer-seller match quality.