Profit-Aware Pricing in Two-Sided Marketplaces: Demand Elasticity, Cost Uncertainty, and Customer Lifetime Value in Brazilian E-Commerce

Dinesh R Poddaturi

Abstract

This paper develops an integrated pricing transformation framework for two-sided e-commerce marketplaces, combining demand elasticity estimation, profit optimization under cost uncertainty, customer propensity modeling, and sequential purchase pattern analysis. Applied to 110,840 transactions across 71 product categories from the Olist Brazilian marketplace (2016 to 2018), the analysis demonstrates that revenue maximization and profit maximization are fundamentally misaligned in elastic markets with material costs: a 40% price cut generates a 160% revenue increase but a 162% profit loss. Profit-optimal pricing derived from the Lerner Index recommends raising Electronics prices by 20% (+5.6% profit), with Electronics accounting for 82% of total optimization opportunity. A key methodological finding is that cost uncertainty dominates elasticity uncertainty: a 5-point COGS error forfeits 70 to 75% of potential profit gains, while a 20% elasticity error preserves profitability across all scenarios. Same-order bundling is non-viable (0% co-purchase rate), but 53% of repeat customers switch categories sequentially, supporting time-triggered recommendation campaigns. A propensity model (AUC 0.577) reveals that freight costs affect retention probability at approximately the same magnitude as product price, validating geographic freight subsidies as a dual acquisition and retention lever. An implementation framework with phased rollout, A/B testing protocols with Bonferroni-corrected stopping rules, and rollback thresholds projects a total annual gain of BRL 56,508 (8.7% above baseline), conditional on cost validation and experimental confirmation.

1 Introduction

1.1 The Revenue Maximization Trap

As algorithmic pricing becomes standard across e-commerce platforms, the tools making these decisions share a common and costly blind spot: they optimize for revenue, not profit. This distinction matters more than most pricing software acknowledges, and the gap between the two is widening as marketplaces scale, cost structures grow more complex, and thin-margin categories become increasingly common. In my experience working across 20+ pricing engagements, the revenue-profit divergence is not an edge case. It is the dominant pattern, and it is getting more expensive to ignore.

Consider an electronics category with price elasticity $\eta = -2.18$ (elastic demand). A revenue-maximizing algorithm recommends a 40% price decrease, projecting a 82.7% revenue increase. At a typical marketplace cost structure of 65% COGS (consistent with industry benchmarks for consumer electronics and tested here across a 40–80% range), that recommendation reduces profit by 144%. The revenue-optimal price falls below marginal cost at current volumes, creating a negative margin spiral where increased sales volume amplifies losses rather than generating profit. The algorithm is technically correct on its own terms and commercially destructive in practice.

The methodology developed here addresses this directly. It draws on applied work across 20+ ongoing client engagements in the U.S., U.K., and Canada, where profit-aware pricing optimization generated $27M in incremental revenue impact over two and a half years. Those engagements covered the full lifecycle of a pricing transformation: estimating demand elasticities from transaction data, integrating cost structures to identify profit-optimal prices, designing and executing A/B tests to validate recommendations causally, managing phased rollouts from pilot to full deployment, and translating findings into executive communication that shifted decision-making from volume metrics to margin metrics. The pattern that emerged across clients and industries was consistent: the analytical tools existed, the data was available, and the revenue opportunity was clear, but the gap between econometric findings and realized business outcomes was almost always a function of implementation discipline rather than analytical precision. Getting the elasticity estimate right mattered less than validating costs first, designing the experiment correctly, and building organizational buy-in for a strategy that would reduce revenue in the short run while improving profit.

Because those engagements involved proprietary client data that cannot be shared publicly, this paper validates the same framework using publicly available Brazilian e-commerce data (Olist, 2016-2018), demonstrating that the methodology transfers across industries and data contexts. The goal is not to report findings about Brazilian e-commerce specifically but to provide a transparent, reproducible example of how an applied economist approaches a complete pricing transformation, from raw transaction data to actionable business strategy, in a setting where every analytical decision can be shown, every assumption can be tested, and every recommendation can be traced back to the data that generated it.

1.2 Why Marketplaces Amplify the Problem

Two-sided marketplaces such as Amazon, eBay, Etsy, and Mercado Livre face structural features that make the revenue-profit divergence particularly acute, and understanding these features is a prerequisite for designing pricing strategy that works in practice rather than in theory.

From the platform perspective, commissions are earned on revenue, not profit, creating incentive misalignment with sellers who care about margins. Platforms cannot directly observe seller costs, limiting their ability to provide profit-aware guidance without access to seller profit and loss (P&L) statements. From the seller perspective, most small and medium businesses optimize for volume and sales rank due to limited pricing sophistication, operating in competitive environments where thin margins make pricing errors costly.

This creates an information asymmetry with strategic value. The platform observes aggregate demand patterns across all sellers and categories; individual sellers see only their own transactions. In the Olist dataset analyzed here, the platform observes 112,650 transactions across 3,095 sellers and 71 product categories. Individual sellers typically see fewer than 100 transactions, which is insufficient for elasticity estimation or marketplace-level inference. The platform's informational advantage creates an opportunity to provide profit-aware pricing guidance that individual sellers, working from their own transaction histories alone, cannot generate independently. This paper operationalizes that opportunity.

Two foundational results from the economics of two-sided markets motivate the approach taken here. Rochet and Tirole (2003) and Rochet and Tirole (2006) establish that platforms must carefully design fee structures across both sides of the market, since price allocation, not just price level, directly affects participation and transaction volume, creating a built-in tension between platform and seller incentives. Hagiu and Wright (2015) extend this, showing that when sellers have superior knowledge of their own products and customers, giving sellers pricing autonomy generates more value for both sides. Together, these results motivate the profit-aware guidance developed in this paper. By helping sellers price for sustainable profitability rather than volume, platforms align their long-term commission revenue with seller success, reducing churn and strengthening marketplace health.

Marketplace structure further shapes strategic opportunities. Analysis of the Olist data reveals extreme seller concentration (Gini coefficient 0.75): the top 20% of sellers generate 82% of revenue. Geographic complexity adds another layer, as São Paulo dominates as a hub-and-spoke distribution center while remote regions face 50-80% freight premiums, with freight costs correlating at $r = -0.576$ with order volume. Customer behavior departs from traditional retail assumptions: 96.9% of orders contain items from a single product category, and only 3.1% of customers return for repeat purchases, a figure substantially lower than the repeat rates typically observed in mature e-commerce platforms, which commonly range from 20 to 30 percent among general merchandise retailers. These structural features invalidate conventional cross-selling and loyalty strategies and require a framework built around the realities of this marketplace rather than assumptions imported from other contexts.

1.3 Research Questions

This paper addresses five practitioner-focused questions spanning pricing optimization, marketplace structure, and customer behavior.

RQ1: When cost structures are integrated into pricing decisions, by how much do profit-optimal prices diverge from revenue-optimal prices, and which categories face the greatest exposure?

RQ2: How sensitive are pricing recommendations to errors in cost assumptions versus errors in elasticity estimates, and what does this imply for where practitioners should focus their analytical effort?

RQ3: What does seller concentration, geographic freight structure, and category loyalty reveal about pricing power and platform strategy?

RQ4: What is the actual same-order co-purchase rate across product categories, and how should platforms redesign cross-sell strategies based on observed rather than assumed customer behavior?

RQ5: How should observational elasticity estimates be translated into implemented pricing strategy, and what experimental validation, phased rollout, and risk mitigation protocols minimize the gap between analytical recommendations and realized business outcomes?

1.4 Preview of Key Findings

All findings are conditional on this dataset, marketplace structure, and cost assumptions. The 65% COGS (Cost Of Goods Sold) base case represents a reasonable central estimate for general merchandise e-commerce and is stress-tested across a 40 to 80% range throughout..

Revenue $\neq$ Profit. Revenue-maximizing prices would reduce profits by 162% for elastic categories at a 65% COGS structure. Profit-optimal strategy requires raising prices in Electronics (+20%), not the 40% decrease suggested by revenue maximization. Watches and Garden Tools are near-optimal at current prices, requiring only minor adjustments.

Cost sensitivity dominates elasticity precision. A 10% error in cost assumptions can reverse a pricing recommendation entirely, from a price increase to a price decrease. A 20% error in elasticity estimates maintains directional consistency. Practitioners should prioritize cost validation over demand model refinement, a counterintuitive but empirically robust finding with direct implications for how pricing teams allocate their analytical resources.

Marketplace structure is highly concentrated. The top 20% of sellers generate 82% of revenue (Gini 0.75), implying that platform retention strategy should focus disproportionately on a small seller segment. Freight costs create substantial geographic demand barriers ($r = -0.576$), with remote regions facing 50-80% premiums relative to Southeast Brazil, a structural opportunity for free shipping threshold strategies.

Customer loyalty is category-specific, not platform-specific. With a 3.1% repeat purchase rate, 93% bucket loyalty within orders, and 53% category switching among returning customers, the data reflects purposeful search-and-purchase behavior rather than platform browsing. This pattern is consistent with the multi-platform aggregator structure of Olist, where customers shop on Mercado Livre or Amazon Brazil without knowing they are purchasing from an Olist seller, limiting the development of platform-specific habits.

Same-order bundling fails; sequential engagement works. The same-order co-purchase rate across focus categories is effectively zero. Customers in this marketplace purchase one category per order, a finding that holds consistently across Electronics, Watches, and Garden Tools and is unlikely to be a data artifact given the 96.9% single-category order rate across the full dataset. Among repeat customers, 53% purchase from a different category in their next order, with a median return time of 29 days. This sequential pattern supports time-triggered recommendation campaigns that replace same-session cross-selling with asynchronous category suggestions timed to the natural return window.

Implementation determines realized value. Quality thresholds ($R^2 \geq 0.62$), A/B testing protocols, phased rollouts, and seller communication frameworks translate analytical findings into realized gains. Without structured implementation, even correctly estimated elasticities and profit-optimal prices generate no business value.

These findings reflect the specific characteristics of this dataset and marketplace context. The 3.1% repeat rate, extreme seller concentration, and Brazilian logistics structure are not universal. US e-commerce platforms will see different numbers. The methodology transfers; the point estimates do not.

1.5 Paper Organization

Section 2 describes the Olist dataset and marketplace structure. Section 3 reviews theoretical foundations across two-sided markets, elasticity estimation, profit optimization, bundling theory, and customer lifetime value. Section 4 details methodology across five analyses: elasticity estimation, profit optimization, bundle analysis, sequential purchase patterns, and CLV modeling. Section 5 presents empirical results by research question. Section 6 provides implementation guidance including A/B testing protocols, phased rollout design, seller communication, and risk mitigation. Section 7 integrates findings into broader business strategy. Section 8 concludes with practitioner recommendations, limitations, and directions for future research.

2 Data and Marketplace Structure

2.1 The Olist Dataset

The analysis uses the Olist Brazilian E-Commerce Dataset, an open-source dataset publicly available on Kaggle.[1] Olist operates as a marketplace aggregator, connecting small and medium-sized businesses with major Brazilian e-commerce platforms. Sellers list products on Olist's centralized platform, which distributes listings simultaneously across Mercado Livre, Amazon Brazil, B2W, and Via Varejo. This aggregator model is analytically important: customers shop on familiar platforms without recognizing Olist as a unified brand, which shapes repeat purchase behavior and limits platform-specific loyalty formation.

The dataset consists of nine interconnected tables spanning September 2016 through August 2018. The analysis period covers January 2017 through August 2018 (20 months), excluding September 2016 as an incomplete launch month and September 2018 due to data cutoff effects. Table 1 summarizes the core dataset dimensions.

[1]Dataset available at: https://www.kaggle.com/datasets/olistbr/brazilian-ecommerce/data. Accessed March 2026.

Table	Records	Content
Orders	99,441	Order lifecycle tracking
Order Items	112,650	Prices, freight costs, seller IDs
Products	32,951	Categories, dimensions, weight
Customers	99,411	Geographic location, unique identifiers
Sellers	3,095	Seller location data
Reviews	99,224	Customer ratings (1-5 stars)
Product Categories	71	Portuguese-English translations
Geolocation	1M+	Coordinates for distance calculations
Order Payments	103,886	Payment methods, installments

Table 1: Olist dataset structure. Nine interconnected tables provide complete transaction coverage across 20 months of stable marketplace activity.

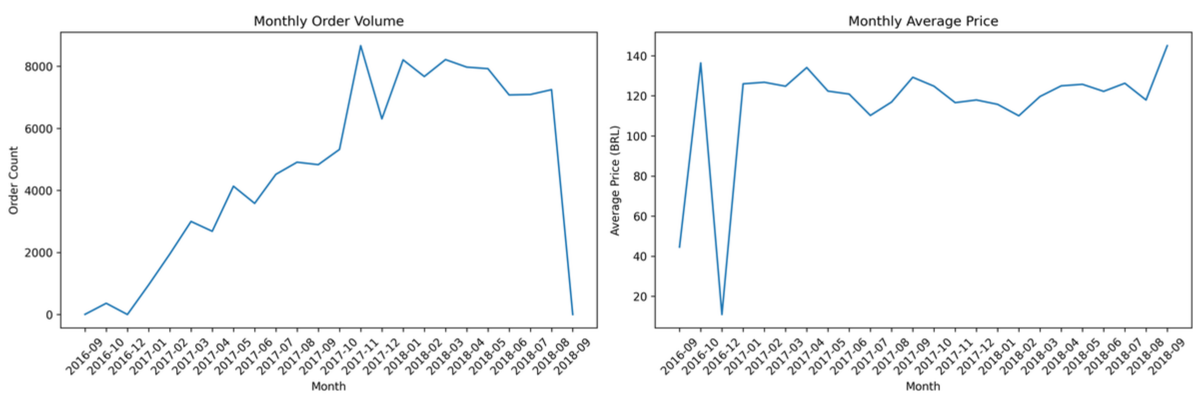

Figure 1: Monthly order volume and average price remain stable throughout 2017–2018, confirming a mature marketplace environment suitable for elasticity estimation.

Monthly order volume grew from near zero in late 2016 to a stable plateau of approximately 8,000 orders per month throughout 2017 and 2018, with average prices remaining relatively stable around BRL 120 (Figure 1). This temporal stability indicates a mature marketplace environment without aggressive promotional cycles or seasonal volatility, which is favorable for elasticity estimation.

The dataset covers all 27 Brazilian states, with primary market concentration in São Paulo (38%), Rio de Janeiro (13%), and Minas Gerais (12%). The average order contains 1.13 items, reflecting predominantly single-item purchases. Review coverage is 99.8% (99,224 of 99,441 orders), providing a near-complete signal of customer satisfaction.

2.2 How Olist Works: The Aggregator Model

Understanding Olist's business model is essential for interpreting the data structure and its limitations. Sellers list products on Olist's platform, pay a commission per sale, and hand items to logistics carriers for fulfillment. Olist pushes these listings to multiple marketplaces simultaneously, so sellers manage inventory centrally while reaching customers across several channels. Marketplaces benefit from access to thousands of sellers without direct onboarding costs, while Olist handles seller quality control and logistics coordination.

From the consumer's perspective, purchases occur entirely within familiar platforms such as Mercado Livre or Amazon Brazil. Consumers do not interact with Olist directly and do not recognize it as a unified brand. This is a critical structural feature: it explains the 3% repeat purchase rate and the absence of platform-specific loyalty, since customers have no awareness of Olist as a destination to return to.

A related data limitation follows directly from this structure. Each transaction in the dataset occurred on one specific marketplace, but platform identity is not recorded. Elasticity estimates therefore represent consumer price sensitivity aggregated across all platforms, not platform-specific demand. This is analogous to grocery demand studies using scanner data aggregated across multiple retail chains, or airline pricing studies aggregating bookings across direct and indirect channels. For seller decision-making, this aggregation is appropriate: sellers using Olist set one price and observe aggregate demand across all distribution channels. The analysis answers the question a seller actually faces: if I change my price, how does total consumer demand respond across all platforms?

2.3 Marketplace Structure

Seller Concentration

Seller competition is highly concentrated. Analysis of active sellers reveals a Gini coefficient of 0.75, indicating extreme winner-take-most dynamics consistent with highly concentrated two-sided marketplaces. The top 20% of sellers generate 82% of total marketplace revenue (BRL 10.9M of BRL 13.3M), while the bottom 50% contribute only 3.3% combined (BRL 439K). The single largest seller captures 1.7% of total marketplace revenue (BRL 227K).

Direct product-level competition is rare: 96.3% of products are sold by a single seller, creating effectively monopolistic pricing conditions for individual SKUs. The 3.7% of products with multiple sellers generate 13.3% of revenue, indicating these competitive products represent high-demand segments. Among multi-seller products, 84% are duopolies with a median price coefficient of variation of 5.8%, suggesting tacit coordination or manufacturer-enforced pricing. Price dispersion increases as seller count rises: products with five or more sellers exhibit a coefficient of variation of 10.8%, indicating that coordination breaks down at higher competition levels.

For platform strategy, this concentration has direct implications. Retention focus and differentiated commission structures should target the top 20% of sellers who account for the majority of marketplace revenue. The same concentration that creates this strategic priority also shapes the analytical approach in Section 5: with 96.3% of products having single sellers, elasticity estimates reflect within-category substitution across products rather than within-product competition across sellers.”

Geographic Structure

The marketplace operates as a São Paulo-centric hub-and-spoke network. Approximately 64% of sellers are based in São Paulo state, serving customers nationally through cross-state shipping. Despite this seller concentration, 63.8% of transactions cross state lines, accounting for 68.7% of total revenue, demonstrating that Olist enables small São Paulo-based businesses to reach a national customer base.

This geographic structure creates systematic pricing disparities. Remote states in the North and Northeast pay 50 to 80% freight premiums over Southeast states, with average freight ranging from BRL 17 in the Southeast to BRL 37 in the North region. Freight cost exhibits a strong negative correlation with order volume (r=-0.576), indicating that shipping costs function as substantial demand barriers in remote markets. Cross-state transactions cost customers approximately 30% more on average (BRL 153 total) compared to same-state purchases (BRL 117 total), with higher product prices contributing alongside freight, suggesting selection effects where only high-value purchases justify cross-state shipping costs.

These geographic patterns have downstream effects on customer retention. The logistic regression model in Section 5 estimates that freight costs affect repeat purchase probability at the similar magnitude as product price (coefficient -0.079 for freight and -0.076 for price), making freight reduction strategies as valuable as price optimization for customer lifetime value.

2.3.1 Product Bucket Construction

The Olist dataset provides 71 granular product categories. While granular categories are useful for individual elasticity estimation, strategic analysis benefits from a higher-level grouping that reflects how customers actually shop and how platforms organize merchandising decisions. I aggregated the 71 categories into 10 strategic product buckets based on purchase motivation similarity, cross-category complementarity, and sufficient transaction volume for statistical analysis.

Table 2 summarizes bucket performance. HOME_ESSENTIALS leads in transaction volume (31% of orders), while LEISURE_LIFESTYLE generates the highest revenue share (29%). SMALL_APPLIANCES is notable for its high average transaction value (BRL 354) despite comprising only 1% of orders, reflecting infrequent but high-ticket purchases. The top four buckets (LEISURE_LIFESTYLE, HOME_ESSENTIALS, PERSONAL_CARE, ELECTRONICS_TECH) account for 85% of marketplace revenue and 66% of transaction volume, making them the natural focus for pricing optimization and platform strategy.

Bucket	Orders	Revenue	Avg Price	% Orders	% Revenue	Categories
LEISURE_LIFESTYLE	26,453	3,831,261	145	24%	29%	16
HOME_ESSENTIALS	33,827	3,301,515	98	31%	25%	13
PERSONAL_CARE	16,193	2,071,139	128	15%	16%	4
ELECTRONICS_TECH	17,262	1,901,030	110	16%	15%	9
AUTO_TOOLS	6,682	994,351	149	6%	8%	9
OFFICE_STATIONERY	4,208	504,904	120	4%	4%	2
SMALL_APPLIANCES	993	351,412	354	1%	3%	3
FASHION_APPAREL	3,734	342,649	92	3%	3%	8
FOOD_BEVERAGE	1,167	67,002	57	1%	1%	3
MISC	504	41,332	82	$<$1%	$<$1%	5

Table 2: Product bucket performance. The top four buckets account for 85% of marketplace revenue. SMALL_APPLIANCES has the highest average transaction value (BRL 354) despite low order volume. Revenue and Avg Price are in BRL.

Full bucket-to-category mappings are provided in Appendix A. The bucket structure is used throughout the analysis to organize elasticity estimation, sequential purchase patterns, and customer lifetime value modeling.

2.4 Customer Behavior Patterns

Repeat Purchase Rate:

The marketplace operates in an extreme low-engagement context. Of 96,096 unique customers, 93,099 (97%) make exactly one purchase and never return. Only 2,997 customers (3%) make repeat purchases, the vast majority making exactly two orders (Figure 2). This 3% repeat rate is substantially below repeat purchase rates observed in destination e-commerce platforms, where platform brand recognition and switching costs create structural retention advantages absent in the aggregator model.

Of these 2,997 repeat customers, 2,801 have confirmed delivered orders and form the analytical sample for propensity modeling and sequential purchase analysis. The remaining 196 repeat purchases involve non-delivered orders and are excluded from behavioral modeling to ensure the post-purchase experience driving retention is fully realized.

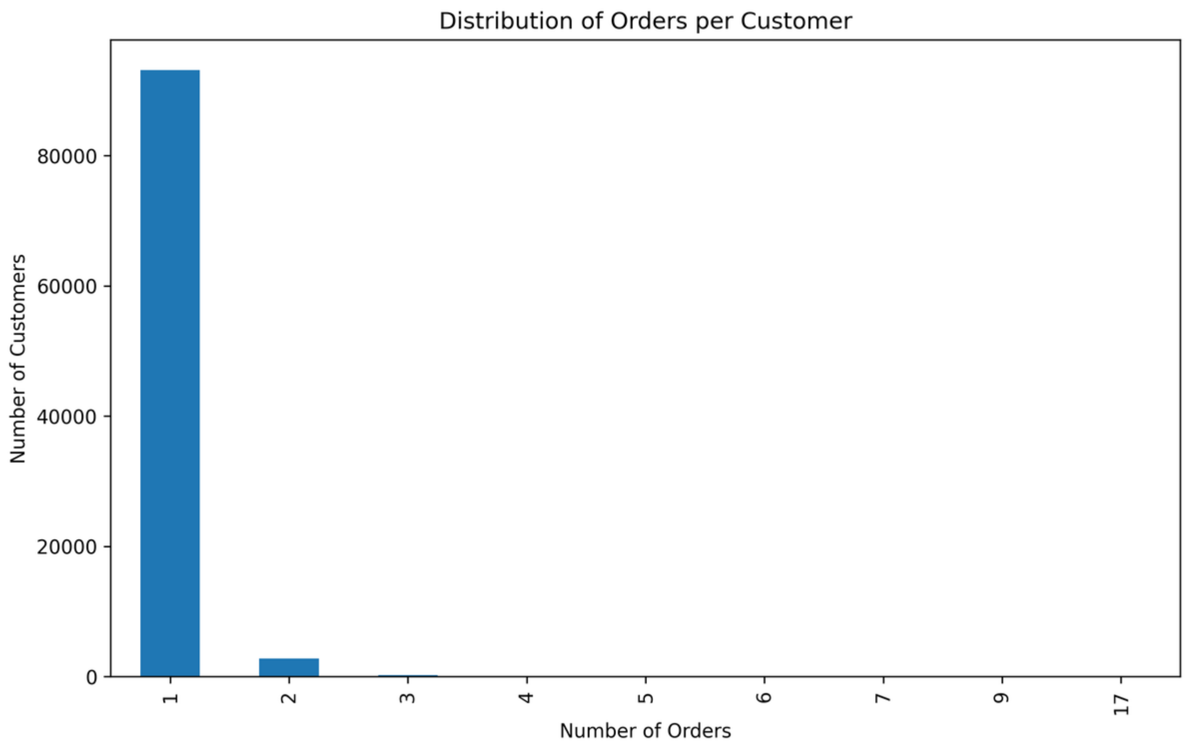

Figure 2: Distribution of orders per customer: 97% of customers make exactly one purchase, with only 3% returning for a second order.

Three structural factors contribute to this pattern. First, the multi-platform aggregator model means customers shop on Mercado Livre or Amazon Brazil without recognizing Olist as a unified destination, so no platform-specific loyalty develops. Second, transaction data shows 98.9% of orders contain items from a single product category, consistent with purposeful search-and-purchase behavior rather than browsing. Third, several high-volume categories (furniture, appliances, tools) have low natural repurchase frequency.

Among the 2,997 repeat customers, the median time between first and second purchase is 29 days, with a mean of 79 days reflecting a right-skewed distribution where some customers return much later. The 75th percentile is 120 days (four months), suggesting that engagement windows for sequential recommendations should extend well beyond the immediate post-purchase period.

Category and Bucket Loyalty

Within individual orders, loyalty is near-total. 98.9% of orders contain items from a single product category, and 93% of orders stay within a single product bucket. Cross-category same-order purchases are effectively zero across focus categories (Electronics, Watches/Gifts, Garden Tools), which has direct implications for bundling strategy discussed in Section 5.

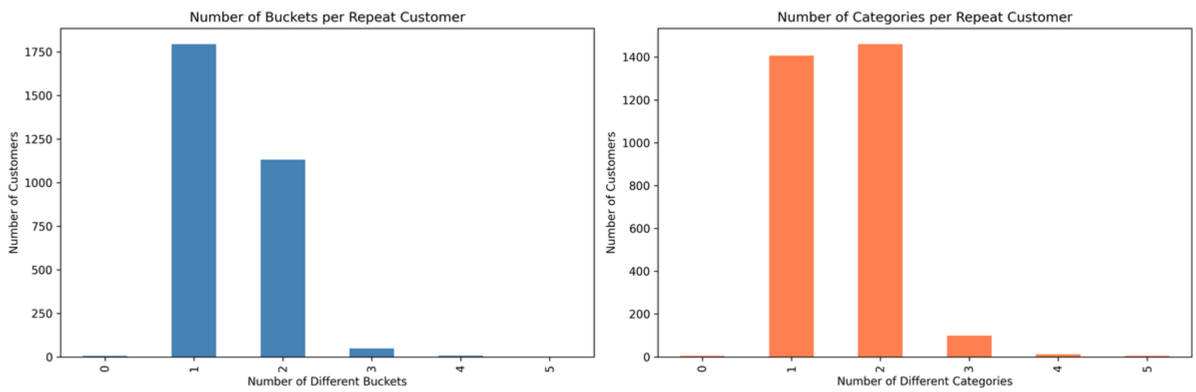

Figure 3: Bucket and category loyalty among repeat customers: most returning customers purchase from only one or two buckets across all their orders.

Behavior across multiple orders tells a different story (Figure 3). Among repeat customers, 60% purchase from only one bucket across all their orders, but 38% switch to a second bucket. At the category level, 47% remain loyal to a single category while 53% switch to a different category in their next order. This 53% category-switching rate among repeat customers creates meaningful opportunity for sequential recommendations, even though same-order bundling is not viable.

The most common sequential transitions are informative for recommendation design. After purchasing Watches/Gifts, customers next buy housewares (4.4%), fashion bags (4.4%), or bed and bath products (3.7%). After Garden Tools, customers move to furniture and decor (15.5%), housewares (7.8%), or sports and leisure (5.8%). After Electronics, customers shift to computers and accessories (9.5%) or furniture and decor (7.1%). These transition patterns form the basis of the time-triggered recommendation framework developed in Section 7.

2.5 Price and Freight Distributions

Product prices are highly right-skewed, with a median of BRL 74.99 and a mean of BRL 120.65, reflecting concentration in lower-priced items with a long tail of high-value electronics, appliances, and furniture extending to BRL 6,735 (Figure 4). Over 76% of transactions fall below BRL 200.

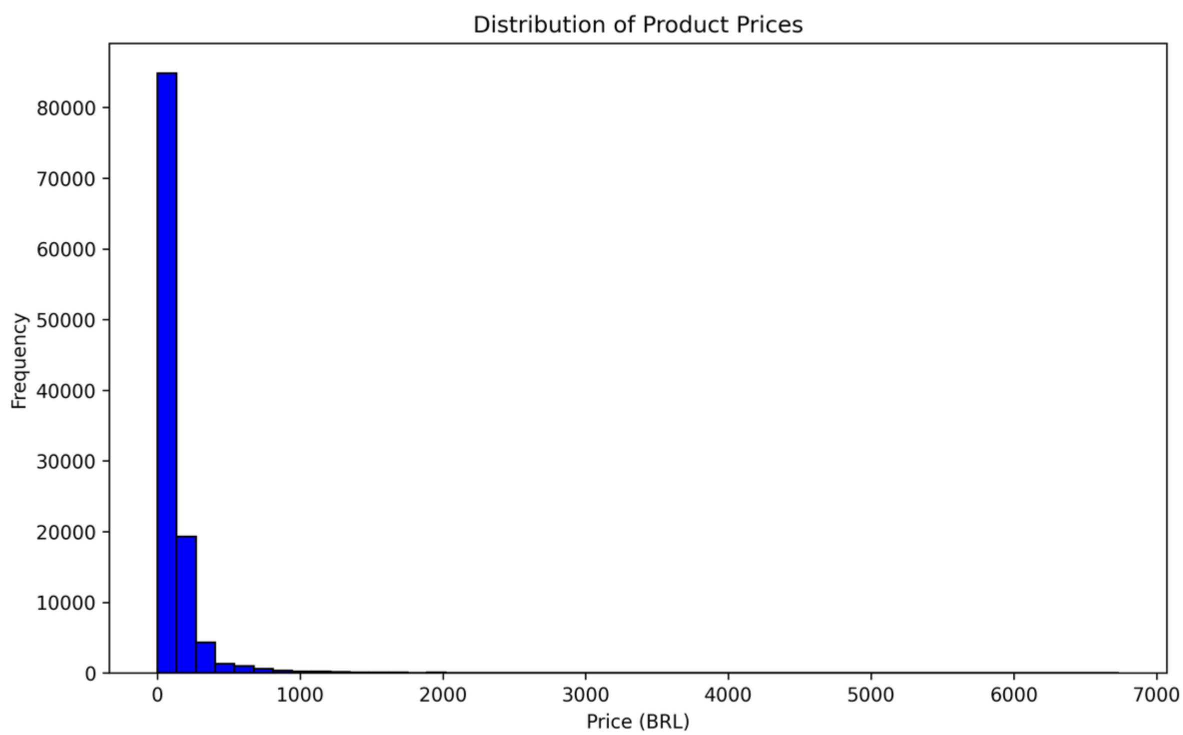

Figure 4: Product price distribution is highly right-skewed with median BRL 74.99 and mean BRL 120.65, reflecting concentration in lower-priced items with a high-value tail.

Freight costs exhibit a banded structure with clear peaks at approximately BRL 15, BRL 30, BRL 45, and BRL 60, reflecting Brazil's standardized shipping zones and carrier weight tiers (Figure 5). The median freight cost is BRL 16.26, representing 18% of the median product price. In remote regions, freight can exceed 50% of product price for low-value items, substantially affecting purchasing decisions and repeat behavior.

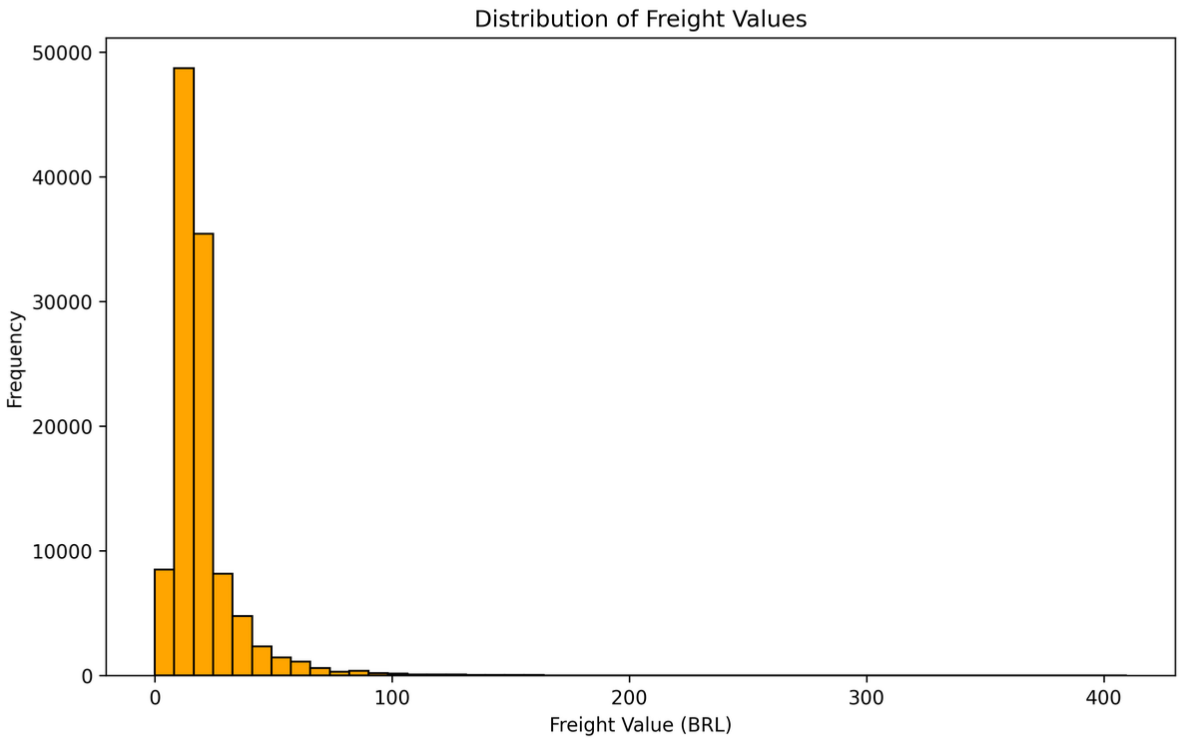

Figure 5: Freight cost distribution shows a banded structure with peaks at BRL 15, 30, 45, and 60, reflecting Brazil's standardized shipping zone and weight tier pricing.

The correlation structure of prices, freight, and weight reveals important pricing dynamics. The moderate positive correlation between price and freight (r=0.41) reflects product heterogeneity: some high-priced items are physically small (watches), while some low-priced items are heavy (garden tools). The stronger correlation between freight and weight (r=0.61) confirms that freight costs are logistics-driven rather than value-driven. The weak price-weight correlation (r=0.34) indicates value-based rather than cost-based pricing across the marketplace.

The total median customer cost is BRL 91.25 (BRL 74.99 product price plus BRL 16.26 freight), with freight representing 17% of the total on average. This freight share rises substantially in remote regions, making geographic pricing strategy a material lever for both demand generation and customer retention.

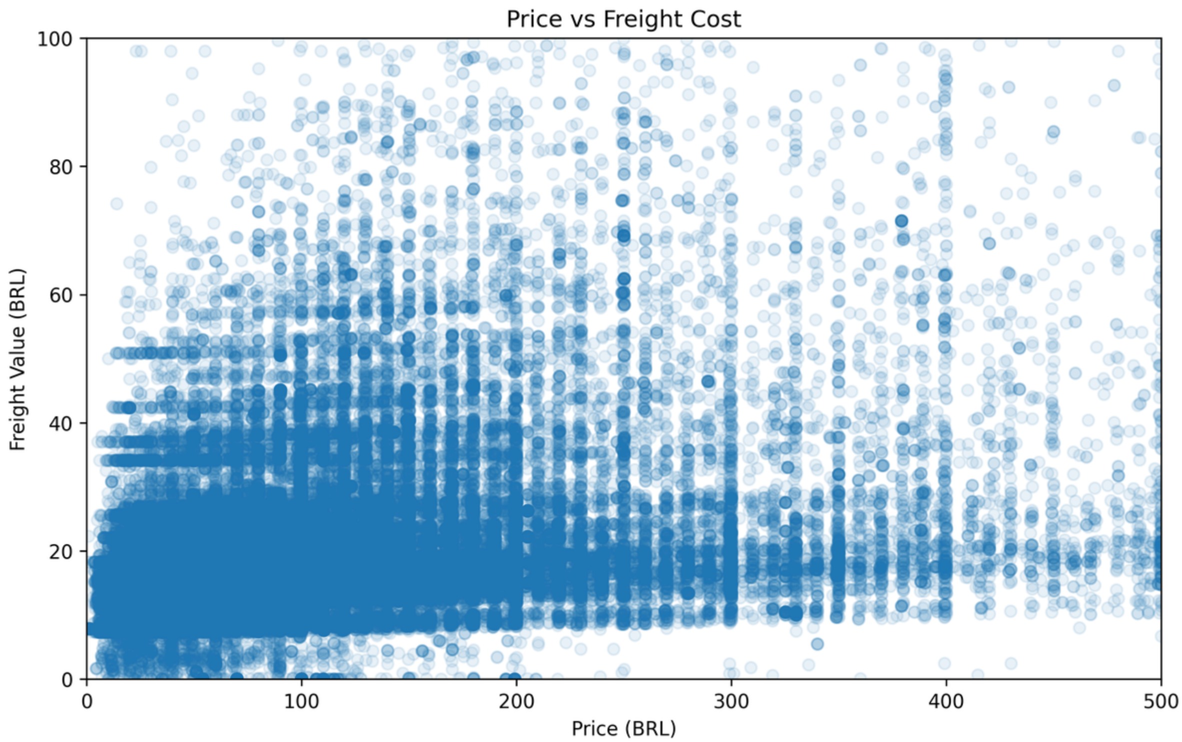

Figure 6: Price versus freight cost shows moderate positive correlation (r = 0.41) with visible banding from standardized freight tiers.

2.6 Data Quality and Limitations

The dataset has several genuine strengths for this analysis. Price, freight, order, product, and seller fields are 100% complete. Review coverage is 99.8%. The 112,650 order items provide adequate statistical power for category-level elasticity estimation. Natural price variation from seller competition and temporal changes creates the dispersion needed for identification. The stable 2017 to 2018 period minimizes confounding from promotional cycles, and all 27 Brazilian states are represented, enabling geographic analysis.

Six limitations bound what can be claimed from this analysis:

Cost structure is unobserved. Seller COGS, fulfillment costs, and inventory carrying costs are not in the data. Profit optimization therefore requires sensitivity analysis across assumed cost scenarios (40 to 80% COGS) rather than precise optimization. Section 5 demonstrates that 10% errors in cost assumptions reverse pricing recommendations, making cost validation more critical than elasticity precision.

Platform identity is not recorded. Which marketplace each transaction occurred on is unknown, so elasticities reflect aggregate demand across platforms. Platform-specific strategies and cross-platform substitution analysis are not feasible with this data.

The time series is limited. Twenty months of data cannot capture long-run elasticities, multi-year seasonal patterns, or structural changes. Results reflect 2017 to 2018 Brazilian marketplace conditions and may not generalize to current market states.

Product-level attributes are sparse. Brand information, product names, and quality descriptors beyond physical dimensions are absent. Category-level aggregation and review score controls partially address this but cannot fully account for within-category heterogeneity.

The repeat purchase sample is small. Sequential purchase patterns are based on 2,997 repeat customers, a consequence of the 3% repeat rate. Findings about sequential behavior should be interpreted as directional patterns rather than precise estimates.

Consumer characteristics are unobserved. Demographics, search behavior, and marketing exposure are not available, limiting controls for consumer heterogeneity. Review scores proxy for quality and category fixed effects absorb time-invariant heterogeneity, but endogeneity from unobserved quality variation may persist.

These six limitations define the boundary conditions of the analysis rather than invalidate it. The analysis can credibly claim category-level price elasticities for robust estimation cases, revenue-profit divergence under validated cost scenarios, marketplace structure patterns specific to the Olist aggregator model, and directional sequential purchase patterns among the repeat customer base. What it cannot claim is equally important: precise profit-optimal prices require cost validation this data cannot provide, findings do not transfer to US e-commerce or current market states without revalidation, and elasticity estimates should not be treated as context-independent.

3 Literature Review

This study sits at the intersection of four research streams: price elasticity estimation in e-commerce, profit-aware pricing versus revenue maximization, customer lifetime value and retention modeling, and two-sided marketplace dynamics. This section positions the paper's contribution relative to prior work, identifies gaps the analysis addresses, and explains how the integrated framework advances both theoretical understanding and practical implementation.

3.1 Price Elasticity Estimation in E-Commerce

3.1.1 Methodological Foundations

The estimation of price elasticity in retail contexts has evolved from early meta-analyses of scanner data to sophisticated approaches leveraging large-scale e-commerce datasets. Tellis (1988) synthesized 367 elasticities from 42 studies across approximately 220 brands and markets, finding a mean price elasticity of -1.76 for consumer goods. Critically, the paper documents substantial variation across product categories, brand life cycle stages, estimation methods, and national settings, cautioning against applying a single benchmark across all contexts.

The shift from brick-and-mortar scanner data to online marketplaces created new opportunities and challenges for elasticity estimation. Chevalier and Goolsbee (2003) pioneered online elasticity estimation by converting Amazon and Barnes and Noble sales rank data into quantity proxies via a Pareto distribution, finding own-price elasticities of -0.45 at Amazon and -3.5 at BN.com. Their panel differencing approach with book fixed effects addressed endogeneity concerns common in observational pricing studies. However, their methodology requires direct price competition between sellers offering identical products, limiting applicability to marketplaces where most products have single sellers, as in the Olist setting analyzed here.

Einav et al. (2015) advanced elasticity estimation by exploiting natural variation in pricing and auction parameters when eBay sellers relist identical products. Using matched listing fixed effects across hundreds of thousands of products, they achieve credible identification by controlling for item and seller heterogeneity. Their auction-based methodology differs fundamentally from fixed-price e-commerce, however, limiting direct applicability to marketplaces where most products have a single seller and no direct price competition.

3.1.2 Aggregation Bias and the Unit of Analysis Problem

A persistent challenge in demand estimation is determining the appropriate level of aggregation. Tellis (1988) explicitly documents this problem, showing that temporal aggregation above the biweekly level introduces a positive bias of 0.83 in price elasticity estimates, and that cross-sectional data alone produces a positive bias of 2.78 relative to time series estimates. These findings establish that the unit of analysis is not a technical detail but a fundamental determinant of whether elasticity estimates are reliable.

This paper demonstrates the same problem in a modern e-commerce context. Bucket-level aggregation across 10 strategic product groups yields null results (β= +0.09, p = 0.648) despite economic theory predicting negative elasticities. Disaggregating to 50 product categories successfully recovers negative elasticities for 65% of categories, with four exhibiting statistical robustness across specifications. This confirms that aggregation level choices fundamentally determine whether demand elasticity can be identified from observational data, a finding with direct implications for how practitioners structure their pricing analyses.

3.1.3 E-Commerce Specific Considerations

Recent work has highlighted factors unique to online pricing. Cavallo (2017) documents that online and offline prices at large multi-channel retailers are identical approximately 72% of the time, with similar frequencies and sizes of price changes, though changes are poorly synchronized across channels. Dinerstein et al. (2018) show that platform search design significantly affects both consumer behavior and seller pricing incentives, with price-prioritizing search algorithms reducing seller margins by roughly 20% for homogeneous products, confirming that even in competitive categories, search frictions generate meaningful pricing power. The analysis here explicitly addresses the role of freight costs by including log(freight) as a covariate, revealing that freight elasticity varies substantially across categories (-6.34 to +7.27), confirming that price and freight affect demand independently rather than as a composite total cost.

3.2 Profit Optimization vs Revenue Maximization

3.2.1 The Lerner Index and Optimal Markup Theory

The foundational theory of profit-optimal pricing derives from Lerner (1934), who established that the optimal markup over marginal cost is inversely related to demand elasticity:

$$\frac{P - MC}{P} = -\frac{1}{\eta}$$

This implies that firms facing elastic demand ($|\eta| > 1$) should maintain relatively low markups, while inelastic demand permits premium pricing. The Lerner Index has been applied extensively in industrial organization to measure market power, though practical implementation requires knowledge of both demand elasticity and marginal costs, a combination rarely available in observational settings.

3.2.2 Revenue vs Profit Misalignment in Elastic Markets

A critical but underappreciated implication of the Lerner rule is that revenue maximization and profit maximization diverge sharply when demand is elastic and costs are non-trivial. For elastic goods ($|\eta| > 1$), revenue increases as price decreases, creating an incentive to cut prices aggressively. When marginal costs are positive, however, volume gains from price cuts come at the expense of margin, and revenue maximization becomes unsustainable as a long-run strategy.

Despite this theoretical clarity, empirical applications of the Lerner framework in e-commerce are scarce. Studies such as Brynjolfsson and Smith (2000) and Baye et al. (2004) focus primarily on price levels and price dispersion rather than profit outcomes, reflecting the broader tendency in e-commerce research to evaluate pricing through a revenue lens. This paper contributes by explicitly calculating the profit destruction from revenue-maximizing pricing: a 40% price cut in elastic categories generates substantial revenue gains but reduces profit by 162%, turning profitable operations into sustained losses. The magnitude of this divergence underscores why revenue-focused optimization tools fail marketplace sellers operating at thin margins.

3.2.3 Cost Uncertainty as Binding Constraint

A novel finding of this analysis is that cost uncertainty dominates elasticity uncertainty in determining optimal pricing. While the academic literature focuses on refining elasticity estimates through better identification strategies (Nevo, 2001; Berry et al., 1995), the sensitivity analysis in Section 5.3 shows that a 5-percentage-point error in COGS assumptions (60% versus 65%) creates larger pricing recommendation errors than a 20% elasticity misestimation. This result has a direct practical implication: practitioners should prioritize cost validation through seller surveys and industry benchmarks over elasticity refinement through expensive A/B testing when implementing profit-aware pricing. This finding inverts the typical emphasis in pricing research and reflects the reality that margin precision matters more than demand precision when margins are thin and demand is elastic.

3.3 Customer Lifetime Value and Retention Modeling

3.3.1 CLV Frameworks in Marketing

The customer lifetime value framework recognizes that optimal pricing must account for repeat purchase behavior, not just immediate transaction profit. Gupta and Lehmann (2005) establish that retention rate directly amplifies profit potential: under reasonable discount rate assumptions, each percentage point improvement in retention generates disproportionately large gains in lifetime value. This relationship highlights that pricing decisions cannot be evaluated on transaction margin alone, since a price that maximizes single-period profit may erode retention and destroy long-run customer value. The empirical CLV calculations in Section 5 operationalize this framework for the Olist marketplace setting.

3.3.2 Propensity Modeling for Retention

Logistic regression propensity models have become standard for predicting customer churn and repeat purchase behavior in non-contractual retail settings. Buckinx and Van den Poel (2005) demonstrate in a fast-moving consumer goods (FMCG) retail context that partial defection can be successfully predicted using behavioral features, with logistic regression achieving AUC of 0.828 on holdout data, and confirm that recency, frequency, and monetary value (RFM) variables are the strongest predictors of future defection. Their model benefits from five months of observed purchase history for each customer.

The propensity model achieves AUC of 0.577 against a 3% base rate, representing modest but statistically meaningful discrimination above chance. Given the severe class imbalance and reliance on first-purchase signals only, individual-level predictions carry substantial uncertainty. The model is best used for directional segmentation rather than precise targeting. The more robust finding is the feature importance result: freight costs affect retention at similar magnitude to product price ($\beta_{\text{price}} = -0.079$, $\beta_{\text{freight}} = -0.076$), a finding that holds regardless of overall model accuracy and implies that shipping cost optimization delivers dual benefits through both demand and retention channels.

3.3.3 Linking Pricing to Lifetime Value

Few studies explicitly connect pricing strategy to CLV implications. Venkatesan and Kumar (2004) empirically validate that firms optimizing customer-level profitability while accounting for retention outperform those optimizing transaction-level margins, demonstrating that CLV-based customer selection generates substantially higher future profits than metrics based on past revenue or cumulative past value alone. Their framework is developed in a B2B context with rich behavioral history, however, and does not address first-purchase pricing decisions under high uncertainty about future retention.

This paper contributes a counter-intuitive finding: low-propensity customers have high first-purchase CLV but represent poor retention investment targets. The bottom 50% of customers by repeat propensity contribute 64% of total CLV due to expensive one-time purchases, which challenges the conventional assumption that high-value customers are automatically the best customers for retention investment. Retention ROI is highest for moderate-value, high-propensity customers in the top 20%, suggesting that resource allocation should prioritize retention probability over first-purchase value.

3.4 Two-Sided Marketplace Dynamics

3.4.1 Platform Competition and Seller Independence

Two-sided marketplace theory (Rochet and Tirole, 2003; Rysman, 2009) emphasizes that platforms must balance pricing across both sides of the market. Platform commission structures affect seller pricing strategies, which in turn affect customer demand, creating interdependencies that complicate pricing optimization. Armstrong (2006) shows that platforms hold monopoly power over access to their buyer base and can extract surplus from sellers through commissions even in the presence of platform competition on the buyer side.

A key assumption in these models is that sellers face competitive pressure either across or within platforms, which disciplines pricing. The Olist marketplace departs substantially from this assumption. Extreme seller concentration (Gini coefficient 0.75) and minimal intra-platform competition (96.3% of products have a single seller) imply that most sellers operate as local monopolists within the platform. This structure amplifies seller pricing power and limits the platform's ability to influence prices through commission adjustments or algorithmic interventions, since there is no competitive mechanism to transmit platform-level pricing pressure to individual sellers.

3.4.2 Geographic Pricing and Shipping Cost Effects

Prior work identifies trust and offline convenience as the primary geographic frictions in online commerce. Hortaçsu et al. (2009) document that same-city trade substantially exceeds what distance alone would predict in online auctions, attributing this to the ability to enforce contracts directly when buyer and seller are geographically close rather than to shipping costs per se. Forman et al. (2009) show that proximity to offline retailers reduces consumers' reliance on online channels and dampens sensitivity to online price discounts, again pointing to convenience and availability rather than freight as the binding constraint.

The setting analyzed here reveals a different dominant mechanism. In a developing-country marketplace spanning a geographically large territory with uneven logistics infrastructure, shipping cost differentials rather than trust or offline competition determine effective market boundaries. The geographic analysis in Appendix B quantifies this: a correlation of $r = -0.576$ between freight costs and order volume indicates that remote regions in North and Northeast Brazil are effectively priced out of the market by shipping costs 50 to 80% above the Southeast baseline. Freight affects retention at similar magnitude to price ($\beta_{\text{freight}} = -0.076$, $\beta_{\text{price}} = -0.079$), implying that geographic subsidy programs such as free shipping thresholds deliver compounding benefits across the customer lifecycle that trust-building or offline competition interventions cannot replicate.

3.4.3 Sequential Purchase and Category Loyalty

Cross-category purchase behavior has been studied primarily in offline grocery contexts. Manchanda et al. (1999) show that complementary category pairs exhibit strong co-incidence in shopping baskets and that cross-category dependencies have meaningful implications for pricing and promotion strategy. Sequential category loyalty and cross-category transition patterns in e-commerce marketplaces remain less examined, as basket-level co-incidence in grocery contexts differs fundamentally from the session-separated purchase patterns observed in online retail.

The transition matrix analysis here reveals extreme category specialization in the Olist marketplace: 93% bucket loyalty and 95% within-bucket category loyalty imply narrow, specialized purchase patterns with minimal cross-category substitution over time. Cross-selling is therefore structurally challenging, but categories can be priced independently without cannibalization concerns. The exception is HOME_ESSENTIALS, where a 6.5% within-bucket switching rate reflects sequential complementarity in which furniture purchases trigger bedding and housewares purchases, creating temporal cross-selling opportunities analogous to those documented by Manchanda et al. (1999) in grocery contexts.

3.5 Gap in Literature and Contribution of This Paper

Despite extensive literatures on elasticity estimation, pricing optimization, CLV modeling, and marketplace dynamics, no prior work integrates all four into a comprehensive pricing transformation framework. Existing studies typically address one component: elasticity papers (Chevalier and Goolsbee, 2003; Einav et al., 2015) estimate demand responses but do not translate findings into profit-optimal pricing; pricing optimization frameworks assume elasticities and costs are known; CLV papers (Gupta and Lehmann, 2005; Venkatesan and Kumar, 2004) model retention but do not link findings to category-level pricing strategies; and marketplace papers (Rochet and Tirole, 2003; Hortaçsu et al., 2009) analyze platform economics but rarely address seller-level pricing optimization.

This paper's contribution is synthesis and integration using a single seller-level dataset that simultaneously supports elasticity estimation, cost modeling, retention analysis, and marketplace context. This approach demonstrates how freight costs operate through both demand and retention channels, how cost uncertainty dominates elasticity uncertainty in determining practical recommendations, how extreme category loyalty enables independent pricing across product groups without cannibalization risk, and how a complete implementation roadmap translates these findings into actionable business strategy. The analysis bridges the gap between academic demand estimation, which focuses on identification and causal inference, and business pricing strategy, which requires actionable recommendations despite data limitations.

The closest existing work is Cachon et al. (2005) on retail assortment planning under consumer search, which combines demand modeling with profit optimization but addresses assortment selection rather than price-level optimization, and does not address CLV, marketplace dynamics, or geographic effects. The framework developed here extends beyond single-firm retail to two-sided marketplaces where seller independence, platform commissions, and cross-state shipping create additional strategic complexity.

3.6 Positioning Within Empirical Industrial Organization

This study also contributes to the empirical industrial organization literature on market structure and pricing. The new empirical IO (Berry et al., 1995; Nevo, 2001) develops structural models of demand and supply to recover primitives such as elasticities, marginal costs, and markups from equilibrium data. BLP-style structural estimation provides counterfactual simulation capabilities but requires strong functional form assumptions and is computationally intensive, which limits its accessibility in practitioner settings where speed and transparency are operational requirements.

The reduced-form approach adopted here is a deliberate methodological choice suited to the business context. Log-log regressions are transparent, fast to estimate, interpretable without advanced econometric training, and robust to the misspecification concerns that can affect structural models when functional form assumptions are uncertain. The trade-off is that counterfactual policy simulations are not available, and causal claims are bounded by the identification strategy. Both limitations are acknowledged explicitly, and alternative specifications are used for validation rather than claiming definitive causal effects. This approach reflects a practical philosophy: use the simplest credible method that answers the business question, and be transparent about what the method can and cannot establish.

4 Methodology

4.1 Empirical Framework Overview

The empirical analysis comprises four interconnected components: demand elasticity estimation via panel regression, profit optimization using the Lerner pricing rule under cost uncertainty, customer lifetime value prediction via binary classification, and sequential purchase pattern analysis through transition probability matrices. Each component addresses a distinct business question, and the components interact: elasticity estimates feed profit optimization, retention propensity scores feed CLV calculations, and both inform the implementation framework in Section 6.

The core identification challenge is the observational nature of marketplace data. Prices are endogenous to unobserved demand shocks, seller strategies, and competitive dynamics. Randomized pricing experiments would provide causal identification, but data constraints require instead a careful econometric design that extracts credible elasticity estimates from observational price variation. The approach relies on fixed effects specifications, robustness checks across aggregation levels, and sensitivity analysis to address endogeneity concerns, with explicit acknowledgment of what can and cannot be claimed from the resulting estimates.

The key variables used throughout the analysis are defined as follows. Q_{ct} denotes quantity sold in category c at time t. P_{ct} denotes average price in category c at time t (BRL). F_{ct} denotes average freight cost. R_{ct}denotes average review score on a 1 to 5 scale. Y_j is a binary repeat purchase indicator for customer j, equal to 1 if the customer makes at least two purchases and 0 otherwise. The panel spans $t = 1, \ldots, 20$ months from January 2017 through August 2018.

4.2 Demand Elasticity Estimation

4.2.1 Econometric Specification

Price elasticity is estimated using a log-log specification, which permits direct interpretation of the price coefficient as an elasticity without additional marginal effect calculations. The general form is:

$$\ln(Q_{ct}) = \beta_0 + \beta_1 \ln(P_{ct}) + \gamma' \mathbf{Z}_{ct} + \text{Fixed Effects} + \varepsilon_{ct}$$

where β_1 is the price elasticity of demand, $\mathbf{Z}_{ct}$ is a vector of controls, and ε_{ct} is the idiosyncratic error term. The log-log specification is standard in the demand literature (Tellis, 1988; Nevo, 2001), accommodates the wide price range in the data (BRL 10 to over BRL 1,000) without scale distortion, and reduces heteroskedasticity through variance compression.

The feasible specification differs by aggregation level due to degrees of freedom constraints, and this difference has substantive consequences for identification.

4.2.2 Bucket-Level Specification

At the bucket level, 10 product groups across 20 months yield 130 observations, providing sufficient degrees of freedom for a two-way fixed effects model:

$$\ln(Q_{bt}) = \beta_0 + \beta_1 \ln(P_{bt}) + \alpha_b + \delta_t + \varepsilon_{bt}$$

where α_b are bucket fixed effects absorbing time-invariant differences in category popularity, and δ_t are time fixed effects absorbing common platform growth trends and seasonal patterns. Standard errors are clustered at the bucket level to account for serial correlation within buckets over time.

4.2.3 Category-Level Specification

At the category level, 50 categories with an average of 19 monthly observations each make time fixed effects infeasible: adding 20 time dummies to a regression with 20 observations creates perfect multicollinearity. Category fixed effects are similarly unavailable because each category is estimated in a separate regression, where a category fixed effect is absorbed by the intercept. Given these constraints, two parsimonious specifications are estimated for each category independently.

The baseline specification regresses log quantity on log price only:

$$\ln(Q_{ct}) = \beta_0 + \beta_1 \ln(P_{ct}) + \varepsilon_{ct}$$

The preferred controlled specification adds freight costs and review scores as covariates:

$$\ln(Q_{ct}) = \beta_0 + \beta_1 \ln(P_{ct}) + \beta_2 \ln(F_{ct}) + \beta_3 R_{ct} + \varepsilon_{ct}$$

where $\ln(F_{ct})$ controls for total customer cost beyond product price, and R_{ct} controls for quality signals that may correlate with both price and demand. The controlled specification is preferred because it absorbs two major confounds without exhausting degrees of freedom. Identification comes from cross-sectional price variation across products and sellers within months rather than from temporal price changes, which is appropriate given the stable pricing environment documented in Section 2.

Robustness is assessed by requiring consistency across both specifications. An elasticity is classified as robust if it is negative and statistically significant at the 5% level in both the simple and controlled models. If sign or significance changes between specifications, the simple model is flagged as confounded by freight or quality effects.

4.2.4 Estimation Procedure

Data are aggregated to the category-month level. Categories are selected from the top five buckets by transaction volume, retaining only those with at least 12 months of observations. This yields a category-level panel of 943 observations across 50 categories. Fifty separate OLS regressions are estimated, one per category, with heteroskedasticity-consistent standard errors.

For each category, elasticity robustness is classified as follows. A robust negative elasticity requires $\hat{\beta}_1 < 0$ and $p < 0.05$ in both specifications. A partial result is negative and significant in one specification only. A null result is statistically indistinguishable from zero in both. A positive result is $\hat{\beta}_1 > 0$ and $p < 0.05$, which may reflect Veblen effects, quality signaling, or residual endogeneity.

Focus categories for profit optimization are selected from the robust negative group using additional criteria: $R^2 > 0.5$ in the controlled specification, at least 15 months of observations, and an elasticity magnitude within the economically plausible range of $-5 < \hat{\beta}_1 < 0$.

4.2.5 Identification Strategy and Threats to Validity

The identification assumption is that conditional on controls, residual price variation is orthogonal to unobserved demand shocks: $E[\varepsilon_{ct} \mid \ln(P_{ct}), \mathbf{Z}_{ct}] = 0$. This conditional independence assumption is weaker than random assignment but stronger than simple correlation. Three threats to this assumption are acknowledged.

The most serious threat is endogenous pricing. Sellers may raise prices in response to positive demand shocks, creating upward bias toward zero or positive elasticities. At the bucket level, the bias is evident: pooled OLS yields $\hat{\beta}_1 = +2.17$, which is reduced to +0.09 after fixed effects. At the category level, freight and review controls absorb some confounds, but category-specific time-varying demand shocks correlated with price changes cannot be ruled out. The consequence is that robust negative elasticities found despite this upward bias are conservative lower bounds on true price sensitivity — actual price sensitivity is likely stronger than the estimates suggest.

A second threat is product mix shifts at the category level. When average category price rises, it may reflect a compositional shift toward higher-priced products rather than a price increase on existing products. This effect biases estimates toward positive elasticity. The category-level specification reduces this problem relative to bucket-level aggregation by narrowing the range of products included in each regression, but within-category quality tiers remain a concern. The R^2 threshold for focus category selection partially addresses this: categories where the model fits poorly are excluded, reducing the influence of noisy composition effects.

A third threat is measurement error in average prices when category price distributions are highly skewed. Quantity-weighted average prices are used throughout to reduce the influence of outlier transactions, but within-category price heterogeneity across product types remains a limitation that brand-level data would resolve.

Instrumental variables estimation is not feasible in this setting. Geographic distance between seller and buyer violates the exclusion restriction because distance affects demand directly through delivery time and local preference. Competitor pricing instruments are unavailable because 96.3% of products have a single seller. Input cost shocks and consumer demographic instruments are not observed in the data. The analysis therefore relies on specification comparisons and robustness checks rather than two-stage least squares for addressing endogeneity.

4.3 Profit Optimization Framework

4.3.1 The Lerner Pricing Rule

Given elasticity estimates $\hat{\eta}_c$ for category c, profit-optimal prices are derived under monopolistic competition with constant elasticity demand. The profit function is:

$$\pi_c(P_c) = (P_c - C_c) \cdot Q_c(P_c)$$

where P_c is the price decision variable, C_c is marginal cost (COGS per unit), and $Q_c(P_c) = Q_0 \left(\frac{P_c}{P_0}\right)^{\eta_c}$ is demand under the constant elasticity assumption implied by the log-log specification. Taking the first-order condition and substituting the demand elasticity $\frac{dQ_c}{dP_c} = \eta_c \frac{Q_c}{P_c}$ yields the Lerner Index:

$$\frac{P_c - C_c}{P_c} = -\frac{1}{|\eta_c|}$$

Solving for the profit-maximizing price:

$$P_c^* = \frac{C_c \cdot |\eta_c|}{|\eta_c| - 1}$$

The markup factor $\frac{|\eta_c|}{|\eta_c|-1}$ decreases in elasticity, confirming that more elastic demand supports lower sustainable markups. At $|\eta| = 2$ the optimal price is twice marginal cost; at $|\eta| = 5$ it is 1.25 times marginal cost.

4.3.2 Cost Structure Assumptions

Seller costs are unobserved in the data. The baseline assumption of 65% COGS (35% gross margin) represents a reasonable central estimate for general merchandise e-commerce, consistent with the mid-range of cost structures observed across the focus categories. Sensitivity analysis tests 60% and 70% to bound recommendations under cost uncertainty, spanning a realistic range for commodity electronics, discretionary gift products, and seasonal garden tools respectively.

Parameterizing cost as a fraction of current price, $C_c = \theta_c \cdot P_{c,\text{current}}$, the optimal price change is:

$$\Delta P_c\% = \left(\frac{\theta_c |\eta_c|}{|\eta_c| - 1} - 1 \right) \times 100$$

For Electronics with $\eta = -2.18$, $\theta = 0.65$, this yields a recommended price increase of 20.1%, derived as $\left(\frac{0.65 \times 2.18}{2.18-1} - 1\right) \times 100$.

4.3.3 Profit Impact and Sensitivity Analysis

After implementing the optimal price P^*, demand responds via constant elasticity:

$$Q^* = Q_{\text{current}} \left(\frac{P^*}{P_{\text{current}}} \right)^{\eta}$$

Profit gain is calculated as $\Delta\pi = (P^* - C) \cdot Q^* - (P_{\text{current}} - C) \cdot Q_{\text{current}}$.

To assess robustness to parameter uncertainty, outcomes are evaluated across a 3×3 sensitivity grid combining three elasticity scenarios (point estimate, 20% more elastic, 20% less elastic) with three cost scenarios (60%, 65%, 70% COGS). This yields nine scenarios per category, identifying whether recommendations are robust across plausible parameter ranges or highly sensitive to specific assumptions.

4.4 Customer Lifetime Value Modeling

4.4.1 Propensity to Repeat Framework

The propensity model asks which first-time customers are most likely to make a second purchase. This is formulated as a binary classification problem where $Y_j = 1$ if customer j makes at least two purchases and $Y_j = 0$ otherwise. The observed base rate is approximately 3%, creating severe class imbalance that constrains model performance. The analysis is restricted to delivered orders, yielding 2,801 repeat customers from the full 2,997, as only completed deliveries generate the post-purchase experience relevant to retention modeling.

The feature vector $\mathbf{X}_j$ includes only information observable at the time of first purchase: log price, log freight, number of items in the order, review score, category indicators for the top 10 categories, and day-of-week and month indicators. This constraint is deliberate: the goal is to identify retention-prone customers from first-transaction signals, which is the operationally relevant prediction problem for acquisition and early engagement targeting.

4.4.2 Logistic Regression Specification

The repeat purchase probability is modeled as:

$$P(Y_j = 1 \mid \mathbf{X}_j) = \Lambda(\beta_0 + \beta'\mathbf{X}_j) = \frac{1}{1 + \exp(-(\beta_0 + \beta'\mathbf{X}_j))}$$

Where $\Lambda(\cdot)$ is the logistic CDF (S-shaped curve mapping $(-\infty, +\infty) \to (0, 1)$).

Log-Odds Formulation:

$$\ln\left(\frac{P(Y_j = 1)}{1 - P(Y_j = 1)}\right) = \beta_0 + \beta'\mathbf{X}_j$$

Coefficients are estimated by maximizing the log-likelihood with an L2 regularization penalty to prevent overfitting given the low base rate. Continuous features are standardized prior to estimation so that coefficients are comparable across variables measured on different scales. The regularization parameter is selected via 5-fold stratified cross-validation, where stratification preserves the 3% repeat rate in each fold.

Model performance is assessed using ROC-AUC, defined as the probability that a randomly selected repeat customer receives a higher predicted propensity score than a randomly selected one-time customer. AUC is preferred over accuracy in this setting because accuracy is misleading under severe class imbalance: a model predicting "never repeats" for all customers achieves 97% accuracy but provides no discriminative value.

4.4.3 Customer Segmentation and CLV Calculation

Out-of-fold propensity scores from cross-validation are used to segment customers into three groups: high propensity (top 20%), medium propensity (20th to 50th percentile), and low propensity (bottom

50%). Each customer's expected CLV is estimated as:

$$\widehat{CLV_j} = V_{j,1} + \hat{P}_j \bar{V}_2$$

where $V_{j,1}$ is observed first-purchase value, $\hat{P}_j$ is predicted repeat propensity, and $\bar{V}_2 =$ BRL 104.48 is the mean second-purchase value estimated from the 2,801 repeat customers with confirmed delivered orders. This is a plug-in estimator that treats the empirical mean second-purchase value as the prediction for all customers' hypothetical second purchase, which is a reasonable approximation given no additional information about individual second-purchase behavior is available.

4.5 Sequential Purchase Analysis

4.5.1 Transition Probability Estimation

Sequential purchase patterns are estimated from the 2,801 repeat customers who generated 3,142 sequential purchase pairs. For each category pair (c, c'), the transition probability is estimated as:

$$\hat{P}(c' \mid c) = \frac{N_{c \to c'}}{N_c}$$

where $N_{c \to c'}$ is the count of transitions from category c to category c' and N_c is the total number of transitions originating from c. The resulting transition matrix is row-stochastic by construction. Diagonal elements represent category loyalty (same-category repurchase) and off-diagonal elements represent category switching. Bucket-level transition matrices are constructed by aggregating category-to-category transitions to the 10-bucket level. The time-to-return distribution is estimated from inter-purchase durations $\Delta t_j = t_{j,2} - t_{j,1}$ for all repeat customers, with quartiles used to define recommendation timing windows.

4.6 Limitations and Robustness

The central limitation of the elasticity analysis is that estimates are conditional correlations rather than causal effects. Endogenous pricing creates upward bias toward zero, meaning robust negative elasticities are conservative lower bounds on true price sensitivity. The absence of valid instruments prevents two-stage least squares correction. Cost uncertainty is the binding constraint for profit optimization: the sensitivity analysis in Section 5.3 demonstrates that a 5-percentage-point error in the COGS assumption generates larger recommendation errors than a 20% elasticity misestimation, making cost validation the critical path for implementation. Aggregation bias is addressed by comparing bucket-level and category-level results, where the null bucket-level result and the successful category-level identification together validate the choice of aggregation level. External validity is limited to two-sided marketplaces with seller autonomy, elastic durable goods categories, and platforms with limited algorithmic price intervention. Results should not be extrapolated to inelastic necessities, subscription models, or B2B e-commerce without revalidation.

All analyses are implemented in Python 3.9 using Statsmodels for regression and Scikit-learn for the propensity model. A fixed random seed ensures reproducibility across all cross-validation procedures.

5 Results

5.1 Elasticity Estimates

5.1.1 Bucket-Level Elasticity: Aggregation Bias

Price elasticity was first estimated at the bucket level using monthly aggregated data (N=130 observations, 10 buckets across 20 months). Table 3 presents results from four nested specifications.

Model	Specification	Elasticity	Std Error	p-value	R^2
1	Pooled OLS	+2.17***	0.35	<0.001	0.24
2	Time FE	+2.06***	0.26	<0.001	0.69
3	Bucket FE	+1.47**	0.56	0.009	0.42
4	**Bucket + Time FE**	**+0.09**	**0.19**	**0.648**	**0.97**

Table 3: *Note: ***p<0.01, **p<0.05. Model 4 (preferred specification) includes both bucket and time fixed effects.*

The preferred specification (Model 4, two-way fixed effects) yields an elasticity of +0.09 (SE=0.19, p=0.648), statistically indistinguishable from zero. The progression from Model 1 (+2.17) to Model 4 (+0.09) reveals the identification problem: pooled OLS captures spurious positive correlation between price and quantity as the marketplace expands over time. Adding time fixed effects controls for platform growth. Adding bucket fixed effects controls for time-invariant category differences. The full specification absorbs both sources of variation, leaving insufficient residual variation for price effect identification.

This null result is informative rather than inconclusive. Bucket-level aggregation masks true price effects through product mix shifts: when bucket-level average price rises, it may reflect customers switching from lower-priced to higher-priced products within the bucket rather than any individual product becoming more expensive. With only 10 buckets and 30 fixed effect parameters in a 130-observation dataset, the approach is overparameterized for the available variation. The appropriate response is disaggregation to category level, not abandonment of the elasticity approach.

5.1.2 Category-Level Elasticity:

Disaggregating to 50 product categories (N=943 observations, approximately 19 months per category) provides sufficient variation for elasticity estimation while maintaining meaningful product differentiation. Table 4 presents estimates from the simple specification and Table 5 from the controlled specification adding log(freight) and review scores.

Rank	Category	Bucket	Elasticity	Std Error	p-value	Sig
1	sports_leisure	LEISURE	-4.05	1.20	<0.001	***
2	watches_gifts	LEISURE	-2.98	0.26	<0.001	***
3	garden_tools	HOME	-2.76	0.39	<0.001	***
4	auto	AUTO	-1.77	1.02	0.10	*
5	toys	LEISURE	-1.57	1.10	0.17	
6	electronics	ELECTRONICS	-1.55	0.34	<0.001	***
7	housewares	HOME	-1.24	1.11	0.28	
8	consoles_games	ELECTRONICS	-1.09	0.23	<0.001	***
9	books_general_interest	LEISURE	-0.96	0.61	0.13	
10	air_conditioning	HOME	-0.71	0.48	0.16	
11	pet_shop	LEISURE	-0.54	0.84	0.53	
12	cool_stuff	LEISURE	-0.19	1.26	0.88	
13	fixed_telephony	ELECTRONICS	-0.14	0.10	0.18	
14	telephony	ELECTRONICS	+0.11	0.81	0.89	
15	computers_accessories	ELECTRONICS	+0.79	1.10	0.48	
16	fashion_shoes	FASHION	+0.85	0.41	0.05	*
17	kitchen_dining	HOME	+1.27	0.45	0.01	**
18	fashion_bags_accessories	FASHION	+2.23	0.29	<0.001	***
19	furniture_decor	HOME	+4.60	0.74	<0.001	***
20	bed_bath_table	HOME	+6.59	1.48	<0.001	***

Table 4: *Note: ***p<0.01, **p <0.05, *p<0.10. Top 20 of 50 categories shown. Negative elasticities: 13 of 20 (65%). Mean elasticity: -0.16. Median: -0.62.*

Sixty-five percent of categories show negative elasticities in the simple model, with five significant at the 1% level. Several categories show unexpected positive elasticities (furniture_decor

+4.60, bed_bath_table +6.59), which may reflect quality signaling effects, product mix shifts toward premium lines, or residual endogeneity where sellers raise prices in response to unobserved demand shocks. These categories warrant product-level investigation before any pricing action.

Category	**Price Elast**	**p-value**	**Freight Elast**	**Review**	**R^2**
watches_gifts	-2.98	<0.001***	-0.28	-0.99	0.89
garden_tools	-2.77	<0.001***	-0.18	-0.04	0.72
electronics	-2.18	<0.001***	+7.27	-0.04	0.62
toys	-1.96	0.12	+2.35	-0.10	0.16
housewares	-1.44	0.03**	+6.51	+2.58	0.82
consoles_games	-1.35	<0.001***	+1.26	-0.16	0.56
pet_shop	-1.14	0.21	+4.00	+0.05	0.20
auto	-0.98	0.21	-1.08	+1.54	0.60
musical_instruments	-0.92	0.38	+2.11	+0.51	0.30
home_construction	-0.74	0.28	+1.92	-1.00	0.36
telephony	-0.72	0.51	+2.27	+1.59	0.56
luggage_accessories	-0.53	0.51	+2.13	-0.95	0.51
computers_accessories	-0.05	0.96	+7.15	+0.56	0.46
cool_stuff	+0.58	0.67	-3.47	+1.55	0.37
sports_leisure	+0.63	0.74	-6.34	+0.00	0.63
home_appliances	+0.75	0.36	-1.78	-0.72	0.17
construction_tools	+1.89	0.06*	-2.80	+0.92	0.42
furniture_decor	+2.06	0.08*	-1.38	+0.91	0.80
fashion_bags	+3.02	<0.001***	-2.99	-0.03	0.77
bed_bath_table	+8.26	<0.001***	+4.36	-2.42	0.67

Table 5: *Note: ***p<0.01, **p<0.05,*p<0.10. Negative price elasticities: 13 of 20 (65%). Significantat p<0.05: 7 categories. Mean: -0.03. Median: -0.73. Average* R^2*: 0.53.*

The positive freight elasticity for Electronics (+7.27) warrants attention. In the Electronics category, higher freight costs are associated with higher quantity sold, which is counterintuitive. The most likely explanation is a composition effect: higher-value electronics (laptops, tablets) both command higher shipping costs due to weight and insurance requirements, and generate higher unit sales in absolute terms relative to lower-value accessories. This is a measurement artifact of category-level aggregation rather than a genuine demand relationship, and underscores that freight coefficients at this level of aggregation should be interpreted cautiously.

The sports_leisure category shows a dramatic sign flip from -4.05 in the simple model to +0.63 in the controlled specification, indicating the simple model was capturing spurious correlation driven by freight or quality variations rather than true price sensitivity. This is precisely the kind of confounding that the controlled specification is designed to detect, and it validates the robustness criteria for focus category selection.

Three categories consistently show positive and statistically significant elasticities across both specifications: bed_bath_table (+6.59 simple, +8.26 controlled, $R^2 = 0.67$), furniture_decor (+4.60 simple, +2.06 controlled, $R^2 = 0.80$), and fashion_bags_accessories (+2.23 simple, +3.02 controlled, $R^2 = 0.77$). Three explanations are plausible and not mutually exclusive. First, genuine Veblen effects: in categories where pre-purchase quality assessment is difficult (furniture aesthetics, fashion accessories), higher prices may signal quality and attract rather than repel demand. Second, product mix shifts: if premium product lines within these categories are growing faster than discount lines, average category price and average category quantity rise together as a composition effect rather than a demand relationship. Third, residual endogeneity: sellers in high-demand periods raise prices and observe higher quantity, producing spurious positive correlation that freight and review controls do not fully absorb. All three categories warrant product-level controlled price experiments before any pricing strategy is implemented. Category-level positive elasticities in these cases are diagnostic signals of measurement limitations, not evidence that raising prices will reliably increase demand.

5.1.3 Robust Findings: Four Categories with Consistent Negative Elasticities

Table 6 identifies categories meeting all three robustness criteria: negative in both specifications, statistically significant at the 5% level in both, and consistent in magnitude.

Category	Bucket	Simple	Sig	Controlled	Sig	R^2	Robust
watches_gifts	LEISURE	**-2.98**	***	**-2.98**	***	**0.89**	**Yes**
garden_tools	HOME	**-2.76**	***	**-2.77**	***	**0.72**	**Yes**
electronics	ELECTRONICS	**-1.55**	***	**-2.18**	***	**0.62**	**Yes**
consoles_games	ELECTRONICS	**-1.09**	***	**-1.35**	***	**0.56**	**Yes**
housewares	HOME	-1.24		-1.44	**	0.82	Partial
toys	LEISURE	-1.57		-1.96		0.16	No (n.s.)
sports_leisure	LEISURE	-4.05	***	+0.63		0.63	No (flip)
auto	AUTO	-1.77	*	-0.98		0.60	No (lost sig.)

Table 6: *Note: Robust requires negative and significant ($p < 0.05$) in both models with consistent magnitude.*

Watches and Gifts ($\eta = -2.98$, $R^2 = 0.89$) shows identical elasticity across both models with exceptionally high model fit, indicating this discretionary gift and personal accessory category exhibits strong price comparison behavior. Garden Tools ($\eta = -2.77$, $R^2 = 0.72$) remains consistent across specifications, reflecting the seasonal and deferrable nature of outdoor purchases: customers delay when prices are too high. Electronics ($\eta = -2.18$, $R^2 = 0.62$) strengthens with controls, suggesting freight costs and review scores were suppressing the true price effect in the simple model and that price-comparison behavior is stronger than the uncontrolled estimate suggests. Consoles and Gaming ($\eta = -1.35$, $R^2 = 0.56$) is moderately elastic despite brand loyalty effects, reflecting competitive substitute availability across gaming platforms.

All four categories exhibit elastic demand ($|\eta| > 1$), which has direct implications for the profit optimization analysis that follows.

5.2 Revenue vs Profit Optimization

5.2.1 The Revenue Maximization Trap

Given elastic demand across all four robust categories, a revenue-maximizing algorithm recommends aggressive price cuts: lower prices increase quantity proportionally more than price decreases, expanding total revenue. Table 7 tests this logic against a 40% price cut scenario at the 65% COGS baseline.

Category	Strategy	Price Change	Revenue Impact	Profit Impact
Watches/Gifts	Revenue-Max	-40%	+175%	-165%
	Profit-Max	-2%	+4.1%	+0.1%
Garden Tools	Revenue-Max	-40%	+147%	-159%
	Profit-Max	+1.5%	-2.6%	+0.1%
Electronics	Revenue-Max	-40%	+82.7%	-144%
	Profit-Max	+20%	-19.4%	+5.6%
Combined	**Revenue-Max**	**-40%**	**+160%**	**-162%**
Total	**Profit-Max**	**Mixed**	**+0.3%**	**+0.6%**

Table 7: *Note: Revenue-Max applies -40% price cut. Profit-Max uses Lerner Index optimal pricing at 65% COGS. Combined profit impact shows revenue maximization converts profit to loss.*

Revenue maximization generates a 162% combined profit loss, converting profitable operations into sustained losses despite a 160% revenue increase. The economic mechanism is straightforward: with elastic demand ($|\eta| > 1$), revenue increases monotonically as price decreases because $\text{Revenue} = P \times Q_0 \left(\frac{P}{P_0}\right)^{\eta}$ grows without bound as $P \to 0$. However, profit accounts for costs, and when marginal cost C is non-trivial, there is an interior profit-maximizing price given by the Lerner rule $P^* = \frac{C \times |\eta|}{|\eta|-1}$. For elastic goods with substantial costs, this profit-maximizing price is significantly above the revenue-maximizing price, prioritizing margin over volume.

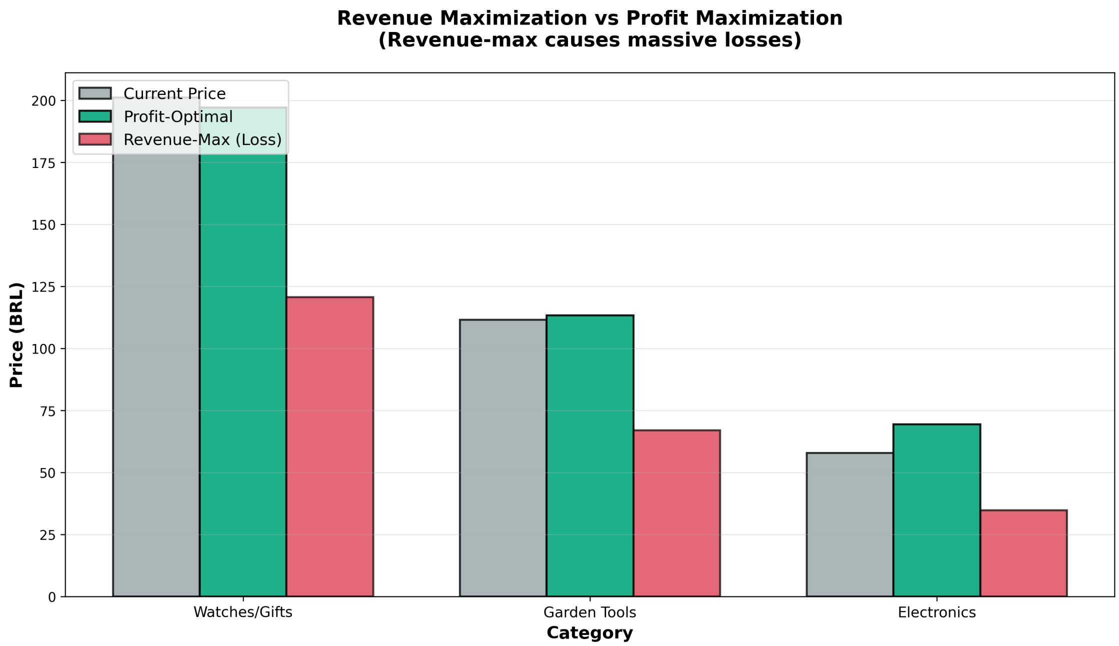

Figure 7: Revenue-maximizing prices (red, –40%) generate large volume gains but destroy profit, while profit-optimal prices (green) range from –2% to +20% and preserve margins.

This finding resolves an apparent paradox in the Electronics recommendation: despite elastic demand ($\eta = -2.18$), the profit-optimal strategy recommends a 20% price increase rather than a price cut. This is not because Electronics demand is inelastic; it is highly elastic. Rather, thin margins at the current price (35% gross margin at 65% COGS) mean that the current price is already below the profit-maximizing level. Volume gained from further price cuts comes at negative marginal contribution. Raising price reduces volume but increases margin per unit more than proportionally, yielding net profit improvement of 5.6% despite a 32.8% volume decline.

5.2.2 Profit-Optimal Price Recommendations

Table 8 presents profit-maximizing prices for the three primary focus categories under the 65% COGS baseline.

Category	Current Price	Optimal Price	Change	Volume Impact	Revenue Impact	Profit Gain
Watches/Gifts	BRL 201	BRL 197	-2.0%	+6.2%	+4.1%	+BRL 575 (+0.1%)
Garden Tools	BRL 112	BRL 113	+1.5%	-4.0%	-2.6%	+BRL 123 (+0.1%)
Electronics	BRL 58	BRL 70	+20.0%	-32.8%	-19.4%	+BRL 3,143 (+5.6%)
TOTAL	–	–	**Mixed**	**-5.4%**	**+0.3%**	**+BRL 3,841 (+0.59%)**

Table 8: *Note: Optimal prices from Lerner Index at 65% COGS. Electronics accounts for 82% of total profit opportunity.*

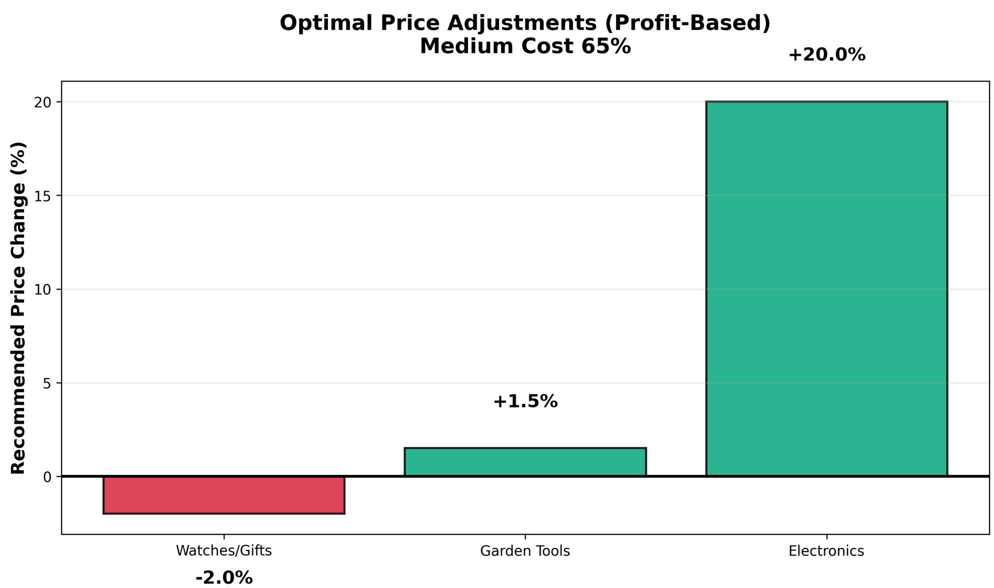

Figure 8: Profit-optimal price adjustments at 65% COGS: Electronics requires a 20% increase, Garden Tools a 1.5% increase, and Watches/Gifts a 2% decrease.

Watches/Gifts and Garden Tools require minimal adjustments (±2%), suggesting sellers have arrived near competitive equilibrium through market experience. The modest profit gains (+0.1% each) confirm these categories are already close to optimal. Electronics is structurally underpriced: the 20% recommended increase reflects thin margins where current pricing does not achieve the optimal Lerner markup. The profit gain of BRL 3,143 represents 82% of total optimization opportunity across all three categories, making Electronics the clear implementation priority despite, and because of, its elastic demand.

The portfolio effect is important: the profit-optimal strategy increases total revenue by 0.3% while increasing profit by 0.59%. Sellers and platforms optimizing on revenue dashboards would reject this strategy. Those optimizing on margin would embrace it.

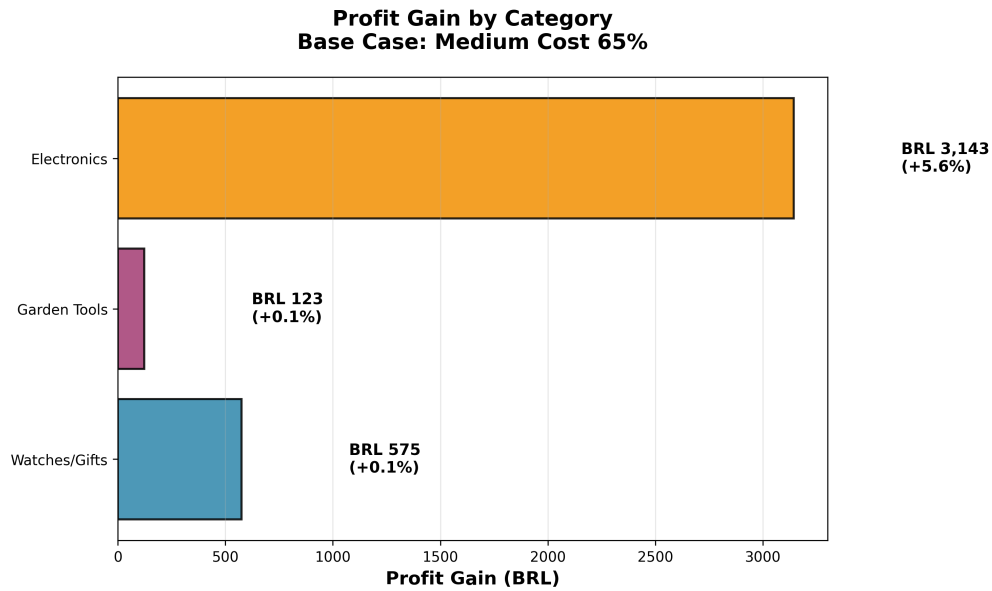

Figure 9: Profit gain by category at 65% COGS baseline: Electronics accounts for BRL 3,143 (82%) of the BRL 3,841 total gain.

5.2.3 Elasticity Sensitivity Analysis

Table 9 evaluates whether recommendations are robust to ±20% elasticity estimation error.

Scenario	**Elasticity Assumption**	**Watches Change**	**Garden Change**	**Electronics Change**	**Total Gain**
Conservative	$\eta \times 0.8$	+12.0%	+18.5%	+40.0%	+BRL 29,646 (+4.6%)
Base Case	$\eta \times 1.0$	-2.0%	+1.5%	+20.0%	+BRL 3,841 (+0.6%)
Aggressive	$\eta \times 1.2$	-10.0%	-7.0%	+5.0%	+BRL 20,773 (+3.2%)

Table 9: *Note: All scenarios at 65% COGS. Parentheses show profit gain percentage. All scenarios generate positive profit gains.*

To test whether recommendations depend critically on elasticity precision, profit outcomes across elasticity scenarios representing ±20% estimation error were evaluated.

All three elasticity scenarios generate positive profit gains, ranging from +0.6% to +4.6%. Elasticity uncertainty affects optimal price levels and profit magnitude substantially, but the key observation is that no elasticity scenario produces a loss. This robustness is reassuring but should not distract from the more important finding in Section 5.3: cost uncertainty creates larger strategic errors than elasticity uncertainty despite generating narrower price ranges.

5.3 Cost Sensitivity Analysis

5.3.1 Why Cost Assumptions Drive Strategy

All recommendations in Section 5.2 assume 65% COGS. Industry benchmarks suggest Electronics COGS ranges from 60 to 75%, Watches and Gifts from 40 to 60%, and Garden Tools from 50 to 70%. Without seller cost data, sensitivity across this range determines whether recommendations are robust or fragile.

Table 10 presents optimal price changes across a 3×3 grid of elasticity scenarios and cost scenarios, spanning the realistic parameter space.

Elasticity Scenario	Cost Scenario	Watches Change	Garden Change	Electronics Change	Total Profit Gain (BRL)	Profit Gain %
Conservative	Low (60%)	+3.5%	+9.5%	+40.0%	10,433	+1.4%
Conservative	Medium (65%)	+12.0%	+18.5%	+40.0%	29,646	+4.6%
Conservative	High (70%)	+20.5%	+27.5%	+40.0%	60,220	+10.8%
Base Case	Low (60%)	-9.5%	-6.0%	+11.0%	15,579	+2.1%
Base Case	Medium (65%)	-2.0%	+1.5%	+20.0%	3,841	+0.6%
Base Case	High (70%)	+5.5%	+9.5%	+29.5%	12,879	+2.3%

Continued on next page

Table 5.8 continued

Elasticity Scenario	**Cost Scenario**	**Watches Change**	**Garden Change**	**Electronics Change**	**Total Profit Gain (BRL)**	**Profit Gain %**
Aggressive	Low (60%)	-16.5%	-14.0%	-3.0%	71,940	+9.7%
Aggressive	Medium (65%)	-10.0%	-7.0%	+5.0%	20,773	+3.2%
Aggressive	High (70%)	-3.0%	+0.0%	+13.5%	3,266	+0.6%

Table 10: *Note: Base Case, Medium Cost row (elasticity × 1.0, 65% COGS) represents central estimates from Section 5.2. Profit gain measured against baseline of BRL 647,678.*

Two patterns dominate. First, cost structure determines strategic direction more decisively than elasticity. Reading down each elasticity column (holding elasticity constant, varying cost from 60% to 70%), recommendations shift by 15 to 18 percentage points per category, with complete directional reversals for Watches/Gifts and Garden Tools. Reading across each cost row (holding cost constant, varying elasticity), recommendations shift by 22 to 35 percentage points but all scenarios within the same cost row remain profitable. Second, a 5-point cost error (65% assumed versus 60% actual, or 65% assumed versus 70% actual) forfeits 70 to 75% of potential profit gains by implementing prices calibrated to the wrong cost structure.

5.3.2 Breakeven Analysis:

Each category has a COGS breakeven ratio at which current pricing is already optimal and zero adjustment is recommended. Table 11 isolates cost sensitivity holding elasticity at base case estimates.

Cost Scenario	**COGS %**	**Watches Change**	**Garden Change**	**Electronics Change**
Low Cost	60%	-9.5%	-6.0%	+11.0%
Medium Cost	65%	-2.0%	+1.5%	+20.0%
High Cost	70%	+5.5%	+9.5%	+29.5%
Range	**10 pts**	**15 pts**	**15.5 pts**	**18.5 pts**
Breakeven	**–**	**66%**	**64%**	**54%**

Table 11: *Note: Breakeven is the COGS percentage where recommended price change equals zero. Electronics breakeven at 54% indicates it is underpriced across the widest realistic cost range.*

Electronics has a breakeven COGS of 54%, meaning that even at unrealistically favorable cost assumptions (46% gross margin), prices should still increase. This makes the Electronics recommendation directionally robust regardless of the precise cost assumption within the plausible industry range of 60 to 75%. Watches/Gifts (breakeven 66%) and Garden Tools (breakeven 64%) are far more sensitive: a 2 to 3 point cost error flips the recommendation from decrease to increase or vice versa.

5.3.3 Cost Error versus Elasticity Error

Table 12 compares the practical consequences of parameter uncertainty across the two dimensions.

Metric	**20% Elasticity Error** ($\eta \times 0.8$ **to** $\eta \times 1.2$)	**10-Point Cost Error (60% to 70% COGS)**
Price swing (Electronics)	35 pts (+5% to +40%)	18.5 pts (+11% to +29.5%)
Price swing (Watches)	22 pts (-10% to +12%)	15 pts (-9.5% to +5.5%)
Price swing (Garden)	25.5 pts (-7% to +18.5%)	15.5 pts (-6% to +9.5%)
Strategic reversals	All 3 categories	All 3 categories
Profit range (optimal)	0.6% to 4.6%	2.1% to 2.3%
Profit foregone if wrong	**Minimal**	**70-75%**
Validation cost	**High (A/B tests)**	**Low (surveys)**
Validation time	**2-3 months**	**2-4 weeks**

Table 12: *Note: Profit foregone calculated as (optimal profit minus achieved profit under wrong assumption) divided by optimal profit.*

Elasticity uncertainty creates wider price recommendation ranges (22 to 35 points versus 15 to 18 points for cost) but preserves profitability across all scenarios: the worst elasticity scenario still generates +0.6% profit. Cost uncertainty creates narrower price ranges but destroys optimization value: implementing 65% COGS prices when true cost is 60% or 70% forfeits 70 to 75% of potential gains. The practical implication is direct: invest 2 to 4 weeks in cost validation via seller surveys before committing to 2 to 3 months of A/B testing for elasticity refinement. Cost validation is the critical path for this analysis, not elasticity precision.

Conditional recommendations pending cost validation are as follows. Electronics should increase prices across the entire plausible COGS range (breakeven 54%, well below industry minimums), making the direction confident even without precise cost data. The specific magnitude should be calibrated once costs are validated: +10 to +15% at 60% COGS, +15 to +20% at 65%, +20 to +25% at 70%, and +25 to +30% at 75%. Watches/Gifts and Garden Tools require cost validation before action: a 3-point cost error reverses direction, and the total profit opportunity (BRL 698 combined) does not justify acting on an uncertain assumption.

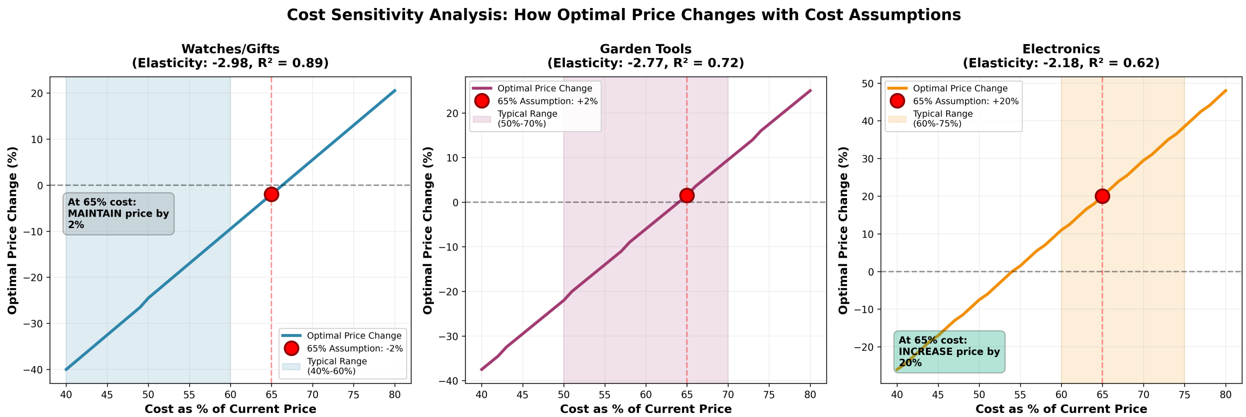

Figure 10: Optimal price change as a function of COGS assumption for each category: Electronics requires price increases across the entire plausible cost range, while Watches/Gifts and Garden Tools cross zero near their 66% and 64% breakeven levels.

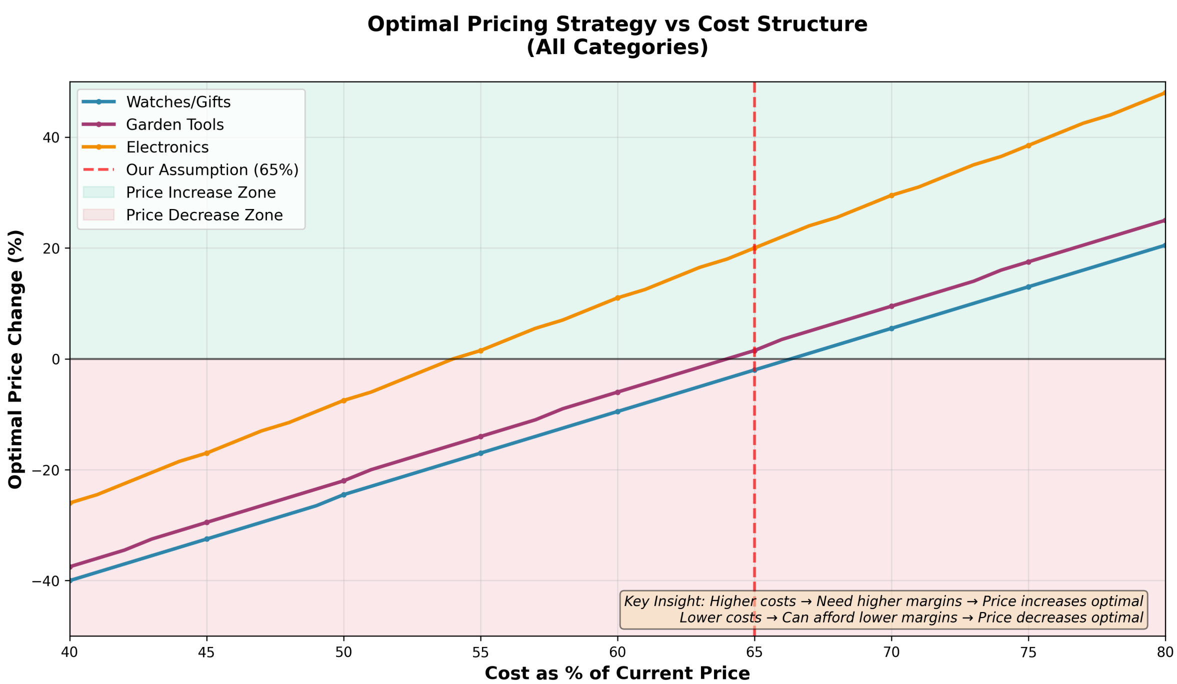

Figure 11: Optimal pricing strategy versus cost structure: higher COGS requires price increases to protect margins, lower COGS enables price reductions to capture volume.

5.4 Bundle Analysis: A Valuable Null Result

Sections 5.2 and 5.3 establish optimal pricing for individual categories. A natural follow-on question is whether cross-category bundling could generate additional revenue by encouraging customers to purchase complementary products together. The answer is no, and the null result is analytically important.

Results: Zero Cross-Category Co-Purchase

Among 11,455 orders containing products from the three focus categories (Watches/Gifts, Garden Tools, Electronics), 98.9% contain products from only one category. The 1.1% multi-category orders (129 orders) do not represent cross-purchases among focus categories.

Category Pair	Co-Purchase Count	Rate
Watches/Gifts + Garden Tools	0	0.0%
Watches/Gifts + Electronics	0	0.0%
Garden Tools + Electronics	0	0.0%
Any Focus Category Pair	**0**	**0.0%**

Table 13: *Note: Zero co-purchases across all three focus category pairs among 11,455 orders analyzed.*

The lift statistic (observed co-purchase rate divided by expected rate under independence) is 0.00x versus an expected 1.5% based on category frequencies. Customers are not merely failing to bundle; they are actively purchasing from single categories with near-perfect consistency.

This behavioral pattern is explained by the marketplace structure documented in Section 2. Customers shop with specific purchase intent: they identify a need, navigate to the relevant category, complete the transaction, and exit. The different buckets occupied by the focus categories (LEISURE_LIFESTYLE for Watches/Gifts, HOME_ESSENTIALS for Garden Tools, ELECTRONICS_TECH for Electronics) occupy separate navigation hierarchies with no natural discovery path between them. There is also no functional complementarity: a customer buying a watch does not simultaneously need a garden shovel.

Platform resources should not be invested in "frequently bought together" widgets, bundle discount engines, or cross-category recommendation algorithms for same-session purchases in these categories. The co-purchase rate of 0% leaves no customers to target. Within-bucket bundling, such as furniture and bedding within HOME_ESSENTIALS, represents a more promising investigation for future work given functional complementarity and shared navigation.

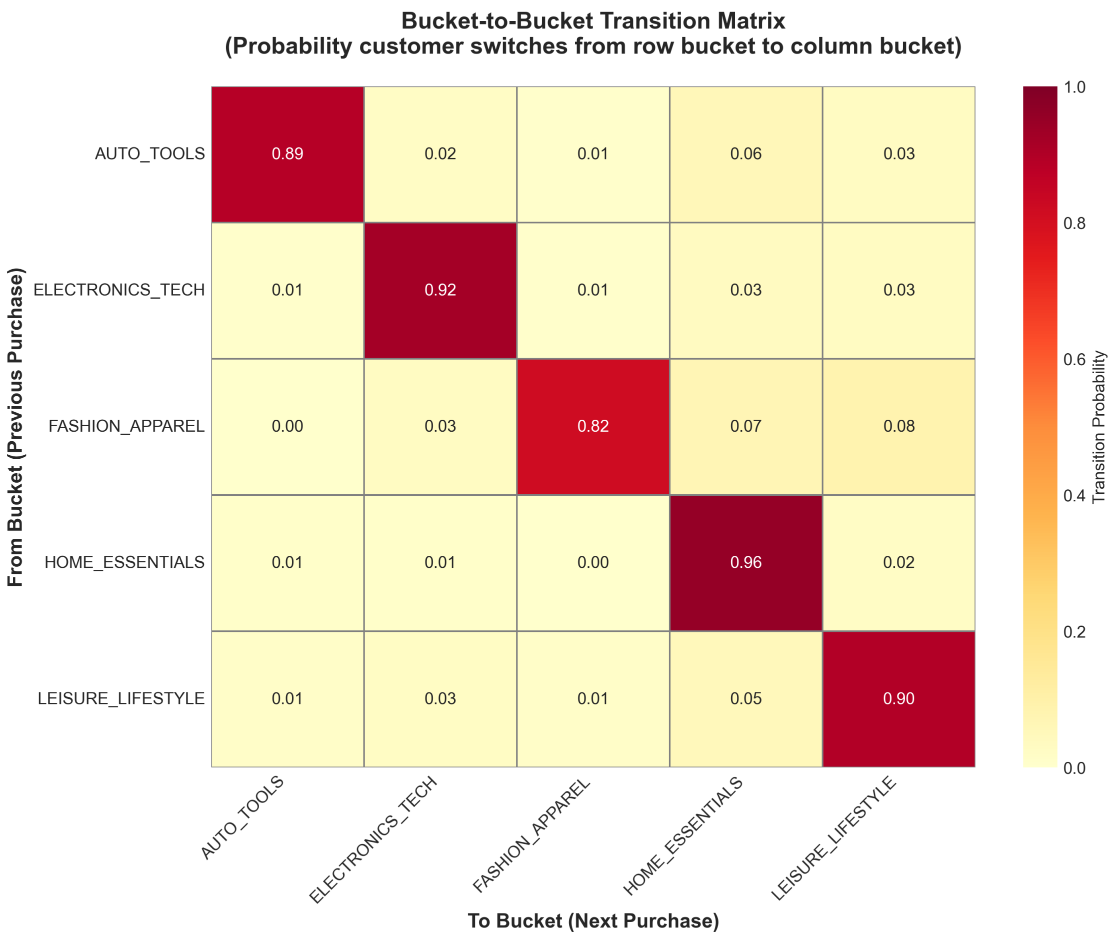

Figure 12: Bucket-to-bucket transition matrix: HOME_ESSENTIALS shows the strongest within-bucket loyalty (96%) and FASHION_APPAREL the most cross-bucket exploration (18%).

This null result prevents wasted implementation effort. Without the co-purchase analysis, a platform might build bundle recommendation infrastructure, launch bundle promotions, and measure performance against unachievable targets. Discovering that bundling is non-viable in this

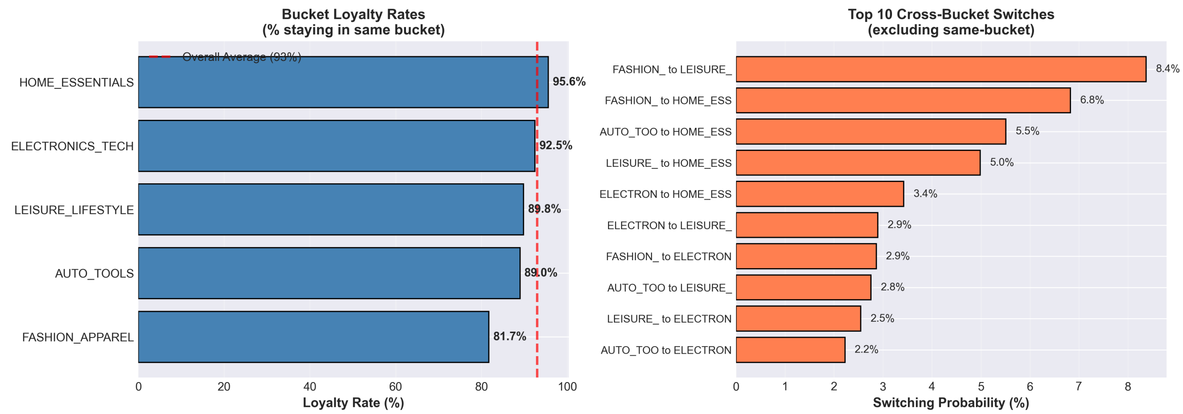

Figure 13: Bucket loyalty rates (left) average 93% across all buckets; top cross-bucket flows (right) concentrate on HOME_ESSENTIALS as destination.

context through two to three weeks of analysis avoids months of failed implementation. Null results are efficient learning, not analytical failures.

The null bundling finding redirects strategic focus from synchronous (same-session) to asynchronous (cross-session) engagement. Section 5.5 demonstrates that while customers do not bundle within orders, 53% of repeat customers purchase from different categories in subsequent orders, creating opportunities for time-based sequential recommendations.

5.5 Sequential Purchase Patterns

5.5.1 Repeat Customer Landscape

Of 93,358 unique customers (filtered to delivered orders), 2,801 (3.0%) make at least two purchases. Among repeaters, 91.9% make exactly two purchases and 6.5% make three purchases, establishing that even among returning customers, engagement is shallow. The median time between first and second purchase is 29 days (mean 79 days, 75th percentile 120 days), reflecting a right-skewed distribution where most customers who return do so within one month, while a long tail returns much later.

5.5.2 Category Transition Analysis

From 2,801 repeat customers (filtered to delivered orders), 3,142 sequential purchase transitions were constructed. Of these, 53.1% involve category switching and 46.9% represent same-category repurchase. This is a striking contrast with within-session behavior (98.9% single-category orders) and reveals two distinct shopping modes: within sessions, customers are intent-driven and narrow; across sessions, they are exploratory and evolving.

From Category	To Category	Count	% of From	Type
bed_bath_table	bed_bath_table	249	58.2%	SAME
sports_leisure	sports_leisure	180	59.4%	SAME
health_beauty	health_beauty	147	61.0%	SAME
computers_accessories	computers_accessories	125	64.1%	SAME
furniture_decor	furniture_decor	124	41.9%	SAME
watches_gifts	watches_gifts	71	53.4%	SAME
home_appliances	home_appliances	59	80.8%	SAME
housewares	housewares	58	35.8%	SAME
fashion_bags_accessories	fashion_bags_accessories	56	49.6%	SAME
furniture_decor	bed_bath_table	46	15.5%	CROSS

Table 14: *Note: Furniture_decor to bed_bath_table (15.5%) is the largest cross-category flow, consistent with sequential room-completion behavior.*

The focus categories exhibit distinct transition profiles. Watches/Gifts shows moderate category loyalty (52.2%) with a median 30-day return window, and cross-category transitions flow primarily to housewares (4.4%), fashion bags (4.4%), and bed and bath products (3.7%). Garden Tools shows the weakest category loyalty (27.2%) with a median 66-day return window consistent with seasonal purchase cycles, and the strongest cross-category flow is to furniture and decor (15.5%), reflecting sequential outdoor improvement projects. Electronics shows very weak loyalty (23.8%)

with a median 52-day return, and transitions disperse widely across categories (45.2% to "other"), suggesting electronics buyers are generalist marketplace explorers rather than category-specific repeat customers.

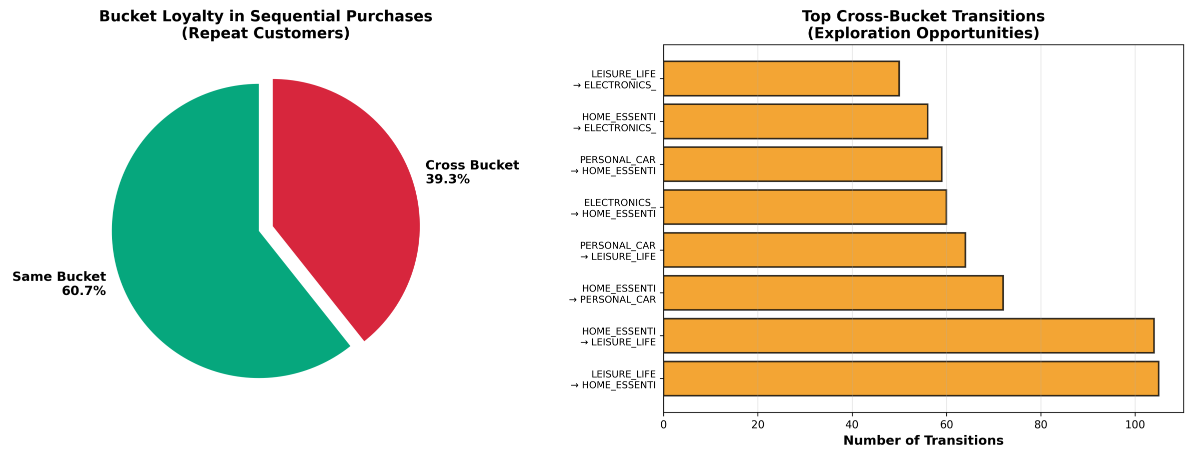

Figure 14: Sequential bucket loyalty (left) and top cross-bucket transitions (right): LEISURE_LIFESTYLE to HOME_ESSENTIALS is the dominant sequential flow.

At the bucket level, sequential loyalty (60.7%) is substantially weaker than within-session loyalty (93.0%), a 32-percentage-point difference that confirms bucket boundaries constrain same-session behavior but become permeable across sessions. The most common cross-bucket flows are bidirectional between LEISURE_LIFESTYLE and HOME_ESSENTIALS (105 and 104 transitions respectively), and between HOME_ESSENTIALS and PERSONAL_CARE (72 and 59 transitions). This cycling pattern, where customers alternate between buckets rather than progressing linearly, supports recurring sequential recommendations rather than one-time cross-sell campaigns.

Table 15 presents the top cross-bucket flows, which are informative for recommendation system design.

From Bucket	To Bucket	Transitions
LEISURE_LIFESTYLE	HOME_ESSENTIALS	105
HOME_ESSENTIALS	LEISURE_LIFESTYLE	104
HOME_ESSENTIALS	PERSONAL_CARE	72
PERSONAL_CARE	LEISURE_LIFESTYLE	64
ELECTRONICS_TECH	HOME_ESSENTIALS	60
PERSONAL_CARE	HOME_ESSENTIALS	59
HOME_ESSENTIALS	ELECTRONICS_TECH	56
LEISURE_LIFESTYLE	ELECTRONICS_TECH	50
LEISURE_LIFESTYLE	PERSONAL_CARE	46
ELECTRONICS_TECH	LEISURE_LIFESTYLE	45

Table 15: *Note: Bidirectional flows dominate (LEISURE $\leftrightarrow$ HOME, HOME $\leftrightarrow$ PERSONAL_CARE), indicating customers cycle through life needs rather than progressing linearly across buckets.*

The dominant pattern is reciprocal rather than directional: LEISURE_LIFESTYLE to HOME_ESSENTIALS and HOME_ESSENTIALS to LEISURE_LIFESTYLE occur at nearly identical rates (105 and 104 transitions). The same bidirectionality appears between HOME_ESSENTIALS and PERSONAL_CARE (72 and 59 transitions). This cycling behavior has a practical implication for recommendation design: after a HOME_ESSENTIALS purchase, recommending LEISURE_LIFESTYLE products is as well-supported as the reverse, and recurring campaigns outperform one-time cross-sell attempts because the flow repeats rather than concluding.

The strategic implication is a direct replacement of "frequently bought together" logic with "customers like you also bought this next month" logic. Time-based triggers at days 21, 29, and 56 (corresponding to the 25th percentile, median, and 75th percentile of the return distribution) capture the majority of returning customers within their natural return window. Category-specific recommendations derived from Table 14 personalize these triggers: Garden Tools buyers receive furniture and decor recommendations; Electronics buyers receive a broader assortment given the high dispersion in their transitions.

5.6 Customer Retention & Propensity Modeling

5.6.1 Model Performance

The logistic regression propensity model achieves a cross-validated ROC-AUC of 0.577 (training AUC 0.583, overfitting gap 0.006). This represents modest but meaningful discrimination above the 0.50 random baseline, achieved under severe constraints: a 3% base rate, first-purchase features only, and no behavioral history beyond the initial transaction. The fold-level AUC ranges from 0.564 to 0.588, indicating stable performance across the data. As discussed in Section 3, a meaningful improvement beyond 0.60 would require behavioral engagement data (browse history, email interactions) and demographic signals unavailable in this dataset. The model's primary value is rank-ordering customers for segmentation rather than precise individual-level prediction.

5.6.2 Feature Importance & Retention Drivers

Table 16 presents the logistic regression coefficients from the full-dataset model.

Feature	Coefficient	Odds Ratio	Interpretation
Positive Drivers:			
cat_bed_bath_table	+0.236	1.27	Bed/bath buyers 27% more likely to repeat
cat_furniture_decor	+0.195	1.22	Furniture buyers 22% more likely
cat_sports_leisure	+0.179	1.20	Sports/leisure buyers 20% more likely
cat_other	+0.157	1.17	Miscellaneous categories 17% more likely
num_items	+0.110	1.11	Each additional item: 11% higher odds
cat_computers_accessories	+0.090	1.09	Computer accessory buyers 9% more likely
cat_health_beauty	+0.090	1.09	Health/beauty buyers 9% more likely
review_score	+0.060	1.06	1-point satisfaction increase: 6% higher odds
Negative Drivers:			
price	-0.079	0.92	BRL 10 price increase: 8% lower odds
freight_value	-0.076	0.93	BRL 10 freight increase: 7% lower odds
month	-0.034	0.97	Slight seasonal decline over time

Table 16: *Note: Coefficients from full-dataset model. Odds ratios >1 increase repeat probability. Price and freight scaled to BRL 10 for interpretability.*

Four findings from this table warrant emphasis. Home essential categories (bed_bath_table +0.236, furniture_decor +0.195) are the strongest positive predictors, reflecting ongoing purchase needs: a customer buying sheets may return for pillows or towels, and a furniture buyer returns to complete a room. This aligns with the sequential transition finding in Section 5.5, where furniture-to-bed-bath is the largest cross-category flow (15.5%). Multi-item first orders predict retention meaningfully: each additional item increases repeat odds by 11%, suggesting that customers who browse multiple products within a session are exploring the platform rather than fulfilling a single specific need, which is associated with higher engagement. Review score has a statistically significant but economically modest effect (+6% per point), indicating that satisfaction is necessary but not sufficient for retention: satisfied customers with one-time needs (wedding gifts, single furniture purchases) do not return regardless of experience quality. Most importantly, price and freight carry nearly similar negative coefficients (-0.079 and -0.076 respectively), establishing that shipping costs affect retention as strongly as product price. This is the empirical foundation for treating geographic freight subsidies as a retention investment rather than purely an acquisition cost.

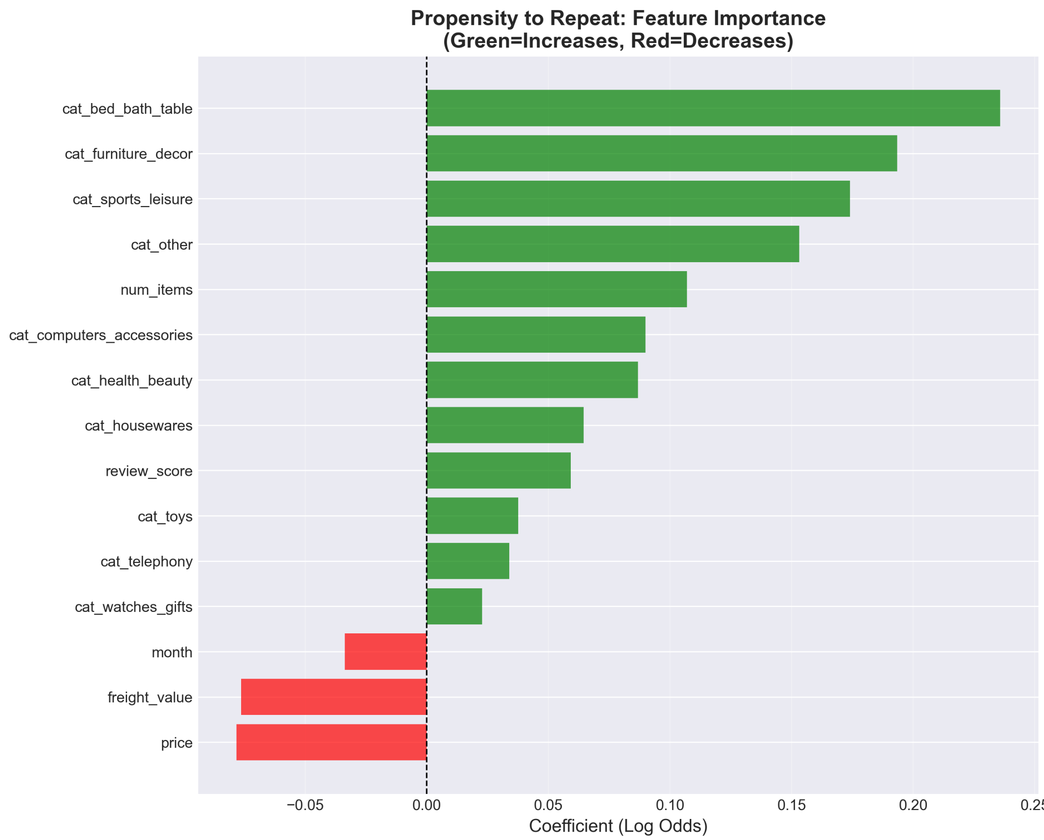

Figure 15: Feature importance from logistic regression propensity model: home category purchases and multi-item baskets increase repeat probability, while higher price and freight reduce it.

5.6.3 Customer Segmentation and CLV

Customers segmented by out-of-fold propensity scores reveal a counter-intuitive CLV structure.

Segment	Customers	Avg Prob	Actual Repeat %	Avg First Purchase	Expected CLV
High (top 20%)	18,671	4.49%	4.40%	BRL 105.72	BRL 110.41
Medium (20-50%)	28,008	3.05%	2.98%	BRL 90.92	BRL 94.10
Low (bottom 50%)	46,679	2.38%	2.45%	BRL 178.20	BRL 180.69
Total	**93,358**	**3.00%**	**3.00%**	–	–

Table 17: *Note: Expected CLV = First Purchase + (Repeat Probability × BRL 104.48 average second purchase value). High segment achieves 4.40% actual repeat rate versus 4.49% predicted, confirming model calibration.*

The low-propensity segment (bottom 50% of customers) contributes 64% of total CLV (BRL 8.43M of BRL 13.13M) through high first-purchase values averaging BRL 178. This inverts the conventional assumption that high-value customers make the best retention targets. Low-propensity customers are expensive one-time purchasers: wedding gifts, single furniture pieces, electronics upgrades. They face structural retention barriers including higher freight (BRL 23.86 average versus BRL 16.84 for high-propensity), lower satisfaction scores (3.92 versus 4.28), and purchase occasions that do not recur. Retention investment in this segment has negative ROI.

The high-propensity segment (top 20%) shows a 4.40% actual repeat rate versus 4.49% predicted, confirming model calibration. These customers buy lower-priced home essentials (BRL 105.72 average) in multi-item orders with positive experiences, and their retention rate can plausibly be lifted from 4.4% to 5.5 to 6.0% through loyalty programs, free shipping incentives, and sequential recommendations timed to the 29-day median return window from Section 5.5. The marginal ROI on retention investment is highest in this segment.

Table 18 presents the behavioral profiles of each segment, which explain the CLV paradox and inform targeted intervention design.

Feature	High (top 20%)	Medium (20-50%)	Low (bottom 50%)
Avg order value (BRL)	105.72	90.92	178.20
Items per order	1.39	1.16	1.03
Avg freight (BRL)	16.84	16.28	23.86
Review score	4.28	4.48	3.92
Actual repeat rate	4.40%	2.98%	2.45%

Table 18: *Note: Low-propensity customers have the highest order values but highest freight costs and lowest satisfaction scores, reflecting structural retention barriers rather than disengagement.*

The contrast between high and low propensity segments is analytically important. High-propensity customers spend less per order (BRL 106 versus BRL 178) but exhibit multi-item purchasing behavior (1.39 items versus 1.03), low freight costs (BRL 17 versus BRL 24), and high satisfaction (4.28 versus 3.92). These are platform-engaged customers buying recurring home essentials, and their repeat behavior is structurally sustainable. Low-propensity customers spend more per transaction but are buying expensive one-time items such as wedding gifts, single furniture pieces, or electronics upgrades, often from remote geographies with high freight. Their low satisfaction and high freight create retention barriers that marketing investment cannot overcome because the purchase occasion itself does not recur. The practical implication is that retention budget should not follow first-purchase revenue: the BRL 178 average order in the low-propensity segment is profitable as a one-time transaction and should be treated as such, while the BRL 106 average order in the high-propensity segment is the foundation for lifetime value programs.

5.6.4 Integration with Pricing and Geography

The propensity model connects the four analytical components of this paper in a single framework. The freight coefficient (-0.076, approximately equal to the price coefficient -0.079) validates the geographic analysis in Appendix B: remote customers paying BRL 37 average freight (versus BRL 17 in the Southeast) face a retention penalty from freight alone equivalent to a substantial product price increase. Free shipping programs therefore generate compounding returns by improving both acquisition (freight correlation with order volume $r = -0.576$, Appendix B) and retention simultaneously. Category coefficients align with the elasticity findings from Section 5.2: bed_bath_table and furniture_decor, which show positive elasticities suggesting quality-signaling dynamics, are also the strongest retention predictors, consistent with premium categories where customers return despite higher prices because ongoing needs exist. Electronics and Watches/Gifts, which are highly elastic and require competitive pricing, show weak or insignificant retention coefficients, consistent with one-time or gift purchase occasions that do not recur. The combined picture is that category portfolio strategy should prioritize home essentials: favorable elasticity-retention profiles where premium pricing is sustainable and repeat behavior is structurally likely.

6 Implementation Framework

Sections 4 and 5 established the empirical foundations: elasticity estimates, profit-optimal pricing recommendations, customer propensity scores, and sequential purchase patterns. This section translates those findings into an actionable implementation strategy, addressing three gaps between analysis and execution. First, cost uncertainty must be resolved before pricing changes are deployed, since a 5-point COGS error forfeits 70 to 75% of potential profit gains. Second, observational elasticity estimates are conditional correlations that require A/B testing for causal confirmation before full-scale rollout. Third, organizational adoption requires phased implementation with clear success metrics and rollback protocols that allow learning without catastrophic downside risk.

The framework is organized into four phases spanning approximately 12 months, followed by detailed guidance on A/B test design, cost validation methodology, sequential recommendation system design, and risk mitigation protocols.

6.1 Phased Rollout Strategy

6.1.1 Phase 1: Cost Validation (Weeks 1 through 4)

The objective of Phase 1 is to narrow COGS uncertainty from the current ±5 point range (60 to 70%) to ±2.5 points (62.5 to 67.5%), which is sufficient to resolve strategic direction for all three focus categories. As Section 5.3 demonstrated, this validation step has higher expected value than any elasticity refinement effort.

Cost validation proceeds through three parallel workstreams. The first is a targeted seller survey of the top 100 sellers per category, asking about COGS percentage, gross margin tier, cost variation by price tier, and primary cost structure challenges. Survey questions should be brief (four questions, under two minutes) and anonymous to encourage honest responses. An incentive of approximately BRL 20 in platform credit per response is sufficient to achieve a 40%+ response rate at a total cost of approximately BRL 6,000 for 300 respondents across three categories. The survey instrument is designed for brevity and honesty. Four questions, completable in under two minutes, cover the essential cost structure information:

Q1	What percentage of your selling price goes to product cost (COGS)? [Slider: 0–100%]
Q2	What are your typical gross margins in this category? (<20% / 20–35% / 35–50% / >50%)
Q3	How do costs vary by price tier? Low-price items (<BRL 50): __% COGS. Mid-price (BRL 50–150): __%. High-price (>BRL 150): __%.
Q4	What is your primary cost structure challenge? (Supplier price fluctuation / Platform commission / Competitor pressure / Other)

Note: Survey administered anonymously with aggregate reporting only. Anonymity increases response honesty for margin-sensitive questions.

The second workstream is industry benchmark research using publicly available sources: IDC and Gartner reports for consumer electronics margins (expected range 60 to 75% COGS), jewelry and gift industry trade publications for Watches/Gifts (expected 40 to 60%), and major home improvement retailer public filings for Garden Tools (expected 50 to 70%). The third workstream, contingent on platform data availability, is a margin back-calculation inferring COGS from observed seller revenues, platform commissions collected, and assumed net margin targets.

If survey results and benchmarks converge within 3 percentage points, proceed to Phase 2 with refined COGS assumptions and updated optimal price calculations. If they diverge by more than 5 points, conduct follow-up interviews with the top 10 sellers per category (30-minute calls) and use the conservative estimate (higher COGS, implying lower recommended price increases) until discrepancies are resolved.

The Phase 1 decision tree governs how validated costs translate into updated recommendations:

Outcome	**Condition**	**Action**
Convergence	Survey and benchmarks within ±3 pts	Proceed to Phase 2 with refined COGS. Update Table 5.6 optimal prices.
Mild divergence	Gap of 3–5 pts	Use midpoint estimate. Flag sensitivity in recommendations.
Strong divergence	Gap >5 pts	Conduct seller interviews (top 10 per category). Use conservative (higher) COGS until resolved.

Note: Conservative COGS assumption (higher cost) implies lower recommended price increases, reducing the risk of over-pricing relative to true cost structure.

As an illustration, if survey and benchmark data converge on 68% COGS for Electronics (versus the 65% base case), the optimal price recommendation updates from +20% to approximately +25%. If Watches/Gifts validates at 55% COGS (versus 65%), the recommendation shifts from -2% to approximately -8%, a directional reversal that confirms the importance of cost validation before any pricing action.

The success metric for Phase 1 is a survey response rate above 40% combined with benchmark validation from at least three independent sources per category.

6.1.2 Phase 2: Category Prioritization and Test Design (Weeks 5 through 8)

Phase 2 selects which categories to test first and designs the A/B experiments. The prioritization framework ranks categories by profit opportunity, elasticity robustness, cost validation confidence, and competitive response risk. Electronics ranks first on all dimensions: it offers BRL 3,143 in projected profit gain (82% of total opportunity), exhibits highly robust elasticity estimates ($p < 0.001$ in both specifications), is structurally underpriced across the full plausible cost range (breakeven COGS 54%), and faces moderate competitive risk in commodity markets. Watches/Gifts ranks second (BRL 575, near-optimal pricing, low competitive risk) and Garden Tools third (BRL 123, near-optimal, seasonal dynamics). The recommended test sequence is Electronics, followed by Watches/Gifts and Garden Tools in parallel once Electronics results are confirmed.

The A/B test for Electronics uses product-level randomization. Product-level randomization is preferred over customer-level or seller-level randomization because it avoids contamination (customers would not see inconsistent prices on the same product), provides sufficient sample size (100 to 200 products per category), and is straightforward to implement. Products are stratified by price tier (low below BRL 40, mid BRL 40 to 80, high above BRL 80) with 50/50 assignment within each tier. The treatment group receives a 20% price increase applied to each product's current price (a product currently priced at BRL 30 moves to BRL 36, one priced at BRL 80 moves to BRL 96), and the control group remains at current prices. This percentage-based treatment preserves the relative pricing structure within the category and tests the directional recommendation from the Lerner analysis without imposing a uniform absolute price that would be inappropriate for products at different points in the price distribution.

The required sample size is derived from a two-sample power calculation targeting detection of a 2% minimum profit increase at 80% power and $\alpha = 0.05$:

$$n = \frac{2\sigma^2(z_{\alpha/2} + z_\beta)^2}{(\mu_1 - \mu_0)^2} = \frac{2(0.40)^2(1.96 + 0.84)^2}{(0.056)^2} \approx 800$$

Where $\sigma = 0.40$ is the estimated standard deviation of profit per product-month expressed as a fraction of baseline profit, reflecting high variance due to product heterogeneity within the Electronics category, and $\mu_1 - \mu_0 = 0.056$ is the expected treatment effect expressed as the same fraction, derived from the +5.6% profit lift projected in Table 8 under the base case assumptions of 65% COGS and $\eta = -2.18$. The $\sigma = 0.40$ assumption is based on typical profit variance observed in consumer electronics retail and should be updated with observed product-level profit variance from the first month of the pilot before finalizing the sample size calculation. If observed variance is substantially higher than 0.40, the required sample size increases proportionally and the test duration should be extended accordingly. If variance is lower, the test achieves higher power than planned, enabling earlier stopping under the pre-specified interim rules.

Using 200 products over 4 months (800 product-months per arm) yields a minimum detectable effect of 4.0%, which is below the expected treatment effect of 5.6% from Table 8, giving approximately 88% power. If the true profit lift is closer to the MDE boundary of 4.0% rather than the projected 5.6%, power drops to 80%, which remains acceptable for decision-making purposes. This design balances statistical rigor with implementation speed, and the interim stopping rules provide additional flexibility to extend the test if early results suggest the true effect is near the lower boundary.

Primary success metrics are profit per product $(P - C) \times Q$ and profit lift relative to control. Secondary metrics include revenue per product and volume per product to verify the elasticity estimate. Guardrail metrics that must not degrade during the test are monitored weekly and trigger investigation or rollback if thresholds are breached:

Metric	Threshold	Action if Violated
Average review score	> -0.10 decline	Investigate quality perception issue
Seller churn rate	>5% increase vs baseline	Survey sellers, consider commission relief
Customer repeat rate	> -0.5pp decline	Check if price shock deterred retention
Competitive undercut rate	>20% of products	Assess competitor response, consider matching

Note: Guardrail violations do not automatically trigger rollback but require investigation within 48 hours. Automatic rollback thresholds are more conservative and detailed in Section 6.5.

6.1.3 Phase 3: A/B Testing Execution and Learning (Months 3 through 6)

The Electronics test launches at the start of Month 3. Daily monitoring tracks revenue, volume, and profit by treatment group, with weekly checks on guardrail metrics and bi-weekly interim analyses for anomaly detection only (not for stopping decisions unless pre-specified thresholds are met).

Interim stopping rules are pre-specified using Bonferroni correction for two interim looks. At Month 2 (50% of data), the test stops only if profit lift exceeds 10% at $p < 0.01$. At Month 3 (75% of data), the threshold is profit lift above 5% at $p < 0.025$. At Month 4 (100% of data), the final decision follows the pre-specified decision rules: a profit lift above 3% at $p < 0.05$ triggers full rollout; a lift of 1 to 3% triggers a 2-month extension; a null result triggers investigation of COGS assumptions and elasticity heterogeneity; and a negative result triggers immediate rollback and model revision.

Outcome	Profit Lift	p-value	Action
Strong success	>3%	$p < 0.05$	Roll out to all Electronics immediately
Moderate success	1–3%	$p < 0.10$	Extend test 2 months, then decide
Null result	-1% to +1%	$p > 0.10$	Investigate: wrong COGS? wrong elasticity? competitor response?
Negative result	< -1%	$p < 0.10$	Immediate rollback, revise model assumptions

Note: Decision rules pre-specified before test launch to prevent post-hoc rationalization of ambiguous results.

After each test, observed results are compared against model predictions and discrepancies are analyzed to update assumptions for subsequent categories. The following example illustrates how this learning loop operates in practice:

Item	Electronics Test (Month 6 Illustrative Results)
Expected profit lift	+5.6% (from Table 8, 65% COGS, $\eta = -2.18$)
Observed profit lift	+16.6% (from A/B test)
Analysis	COGS validated at 68% (not 65%), implying larger margin opportunity. Actual elasticity closer to -1.8 (not -2.18), meaning less volume loss than predicted. Combined effect: higher margins and lower volume sensitivity produce better profit outcome than model projected.
Update for Watches/Gifts	Use validated COGS from survey rather than 65% default. Apply conservative elasticity ($\eta \times 0.9$) for profit projections. Expect model predictions to be conservative if the pattern holds.

Note: This example is illustrative. Actual test results will differ. The learning loop process applies regardless of whether outcomes exceed or fall short of model predictions.

This learning loop is the mechanism by which the framework improves over successive implementation cycles. Each test generates updated parameter estimates that calibrate subsequent recommendations, reducing reliance on industry benchmark assumptions as proprietary cost and elasticity data accumulate.

Three outcome scenarios inform next steps. If the test confirms strong positive results ($lift \geq 3\%$, $p < 0.05$, no guardrail violations), rollout proceeds to all Electronics products in Month 7 and Watches/Gifts and Garden Tools tests launch in parallel in Months 7 through 10. If results are null or weak, the investigation checklist examines whether the COGS assumption was wrong, whether elasticity is heterogeneous across price tiers, whether competitors responded, and whether seasonality confounded the estimate. If results are negative, prices roll back immediately and model assumptions are revised before any further testing.

Watches/Gifts and Garden Tools tests follow in Months 7 through 10. Both categories expect minimal profit gains (0.1% each) and are expected to confirm near-optimal current pricing rather than generate material optimization opportunity. Given that Electronics captures 82% of total projected profit, these tests are confirmatory rather than transformative and can be deprioritized if resource constraints require sequencing.

6.1.4 Phase 4: Full Implementation and Scaling (Month 7 onward)

Upon Electronics test confirmation, all Electronics products move to the validated optimal price in Month 7. The projected annual profit gain from pricing optimization is BRL 37,716 (BRL 3,143 per month for Electronics plus BRL 123 per month for Garden Tools once confirmed).

Retention program rollout begins in Month 8 across three components. A loyalty points program targeting high and medium propensity segments (top 50% of customers by propensity score) offers tiered discounts at 2, 5, and 10 purchases. A free shipping threshold program differentiates by propensity segment, lowering the threshold for high-propensity customers to reduce the freight barrier that the propensity model identifies as equally damaging to retention as product price. A sequential email recommendation program deploys time-based triggers at days 21, 29, and 60 post-purchase, using the category transition probabilities from Table 14 to personalize recommendations.

These are illustrative projections based on industry-average email conversion rates (25% open rate, 15% click-through rate, 5% purchase conversion) and the freight propensity coefficient from Section 5.6. Actual results will vary and each program component requires separate A/B test validation before full deployment. The sequential email component delivers the strongest projected return (255% ROI, BRL 12,762 net gain) at minimal marginal cost, making it the natural first implementation priority. The loyalty points program delivers the weakest projected return (4.3% ROI, BRL 2,093 net gain) because broad enrollment costs are distributed across all participants regardless of actual repeat behavior. A refined version targeting only high-propensity customers (top 20% by propensity score from Section 5.6) would substantially improve loyalty program ROI by concentrating costs on customers most likely to respond. The free shipping threshold program sits between the two (22% ROI, BRL 2,461 net gain) and is particularly compelling because the freight coefficient in the propensity model ($\beta_{\text{freight}} = -0.076$) provides empirical grounding for the retention lift assumption that the email and loyalty projections lack.

Table 19 presents illustrative projected returns for each retention program component, based on industry-average email conversion rates and the freight propensity coefficients from Section 5.6. These are planning estimates rather than predictions and should be treated as lower bounds pending A/B test validation of each component.

Program	**Incremental Revenue**	**Cost**	**Net Gain**	**ROI**
Loyalty Points	BRL 48,772	BRL 46,679	BRL 2,093	4.3%
Free Shipping Thresholds	BRL 13,664	BRL 11,203	BRL 2,461	22.0%
Sequential Emails	BRL 17,762	BRL 5,000	BRL 12,762	255%
Total	**BRL 80,198**	**BRL 62,882**	**BRL 17,316**	**27.5%**

Table 19: *Note: Projections assume 25% email open rate, 15% click-through rate, and 5% conversion rate consistent with industry benchmarks for e-commerce retention campaigns. Actual results require A/B test validation before full deployment.*

Sequential emails generate the highest ROI (255%) because the marginal cost of email delivery is low relative to the incremental purchase value. The loyalty points program generates the lowest ROI (4.3%) because broad program costs are distributed across all enrolled customers regardless of their actual repeat behavior. A refined version targeting only high-propensity customers (top 20% by propensity score) would substantially improve the loyalty program ROI by concentrating costs on customers most likely to respond.

The combined pricing and retention transformation projects the following annual impact by Month 12:

Component	Monthly Impact	Annual Impact
Electronics pricing optimization	BRL 3,143	BRL 37,716
Garden Tools pricing optimization	BRL 123	BRL 1,476
Watches/Gifts (near-optimal, no change)	BRL 0	BRL 0
Pricing subtotal	BRL 3,266	BRL 39,192
Loyalty points program (net)	–	BRL 2,093
Free shipping thresholds (net)	–	BRL 2,461
Sequential email campaigns (net)	–	BRL 12,762
Retention subtotal (net)	–	BRL 17,316
Total incremental profit	–	**BRL 56,508 (+8.7%)**

Table 20: *Note: All figures conditional on cost validation confirming 65% COGS baseline, A/B test confirmation of Electronics elasticity estimate, and retention conversion rates consistent with industry benchmarks. Electronics pricing optimization accounts for 67% of total projected gain.*

By Month 12, the combined pricing and retention transformation projects a total incremental annual profit of approximately BRL 56,508 above the current baseline of BRL 647,678, representing an 8.7% improvement. This figure is conditional on cost validation confirming the 65% COGS baseline, A/B test confirmation of the Electronics elasticity estimate, and retention program conversion rates in line with industry benchmarks. Electronics pricing optimization accounts for 67% of the total projected gain (BRL 37,716 of BRL 56,508), reinforcing the prioritization of Electronics in the phased rollout sequence.

6.2 A/B Testing Design

6.2.1 Statistical Framework

The A/B test estimates the causal effect of a price increase on category-level profitability using a pre-registered design with stratified randomization, pre-specified decision rules, and controlled interim analyses. The intent-to-treat estimator is:

$$\hat{\Delta} = \bar{\pi}_{\text{treatment}} - \bar{\pi}_{\text{control}}$$

For precision improvement, a regression-adjusted estimator controls for price tier:

$$\pi_i = \beta_0 + \beta_1 \text{Treatment}_i + \beta_2 \text{Price Tier}_i + \varepsilon_i$$

with standard errors clustered at the product level to account for within-product correlation across months.

The minimum detectable effect formula establishes test duration requirements:

$$\text{MDE} = \frac{2.8\sigma}{\sqrt{n}} = \frac{2.8 \times 0.40}{\sqrt{800}} = 4.0\%$$

Since the expected treatment effect (5.6%) exceeds the MDE (4.0%), the design achieves approximately 88% power. The Bonferroni-corrected interim stopping rules pay a small power penalty (from 88% to approximately 85%) in exchange for the ability to stop early on strong results or damaging outcomes.

6.2.2 Heterogeneous Treatment Effects

Pre-specified subgroup analyses examine three dimensions of treatment effect heterogeneity. Product price tier (low below BRL 40, mid BRL 40 to 80, high above BRL 80) tests whether low-tier products are more elastic than high-tier products, which would support tier-specific rather than uniform pricing. Seller tier (top 20% versus bottom 80% by revenue) tests whether established sellers with stronger reputations can sustain higher prices with less volume loss. Review score (above 4.5 versus below 4.0) tests whether high-rated products can charge a quality premium without demand loss.

All subgroup analyses are pre-specified and use Bonferroni correction ($\alpha = 0.05/3 = 0.017$)

to control for multiple comparisons. These analyses are exploratory and inform future pricing strategy rather than driving the primary rollout decision.

6.3 Sequential Recommendation System

The sequential recommendation system operationalizes the Section 5.5 findings using three time-based triggers. At day 21 (one week before the median 29-day return window), a personalized email presents the top three sequential recommendations derived from the category transition matrix. At day 28 (at the median return window), a follow-up reaches customers who did not engage with the day 21 email, with urgency framing and a short-validity discount code. At day 60, a re-engagement campaign targets customers who have not returned, offering a broader assortment recommendation and a higher discount to reactivate lapsed buyers.

Trigger exclusions prevent over-communication: customers who have already made a second purchase receive no further retention emails; low-propensity customers whose first purchase was a high-value one-time item (above BRL 300) are excluded from retention campaigns given the negative ROI demonstrated in Section 5.6.

Category-specific recommendation logic is derived directly from the transition probability matrix in Table \ref{tab:cat_transitions:

First Purchase	**Recommend 1**	**Prob**	**Recommend 2**	**Prob**	**Recommend 3**	**Prob**
Furniture	Bed/Bath	15.5%	Housewares	7.8%	Garden Tools	6.5%
Bed/Bath	Bed/Bath (same)	58.2%	Furniture	3.2%	Housewares	2.8%
Watches/Gifts	Watches (same)	52.2%	Housewares	4.4%	Fashion Bags	4.4%
Garden Tools	Furniture	15.5%	Housewares	7.8%	Sports/Leisure	5.8%
Electronics	Computers/Accessories	9.5%	Furniture	7.1%	Electronics (same)	23.8%

Note: Transition probabilities from Section 5.5. Same-category repurchase is listed where it is the highest-probability transition (Bed/Bath, Watches/Gifts, Electronics). For Electronics, same-category repurchase (23.8%) ranks third behind computers/accessories and furniture.

A separate A/B test for email content compares category-specific recommendations (Variant A, driven by transition probabilities) against bestseller recommendations across all categories (Variant B, driven by platform-wide popularity). Click-through rate and conversion rate are the primary metrics. The hypothesis is that Variant A outperforms due to relevance, but empirical validation is required before committing to the category-specific approach at scale.

6.4 Risk Mitigation

6.4.1 Competitive Response

If competitors respond to price increases by undercutting across more than 20% of the Electronics portfolio, the response protocol escalates in stages. Mild responses (10 to 20% of products undercut) require monitoring only. Moderate responses (20 to 35% undercut) trigger selective price matching on the top 10 products by volume, preserving higher prices on long-tail SKUs. Severe responses (above 35% undercut) trigger a pricing committee review with the option of full rollback. In a sustained price war scenario, the non-price competitive levers (faster shipping, higher review scores, loyalty program enrollment, and free shipping thresholds) become the primary response, since competing purely on price in elastic categories erodes margins for all participants.

6.4.2 Margin Floor Enforcement

A hard margin floor of 25% gross margin (equivalent to 75% COGS) is enforced on all pricing decisions. Any proposed price change that would imply a margin below this threshold is rejected automatically and escalated for manual review. The floor can be overridden by a pricing committee for loss-leader strategies, competitive defense on high-volume SKUs, or inventory liquidation, but all overrides are logged and reviewed quarterly. This prevents the revenue-maximization trap identified in Section 5.2 from recurring through gradual margin erosion rather than a single large price cut.

6.4.3 Rollback Protocols

Automatic rollback triggers that require no committee approval are a profit decline exceeding 5% for two consecutive weeks, a seller churn rate increase above 10% versus baseline, a review score decline exceeding 0.2 points versus baseline, and a competitive undercut rate above 50% of products. Manual review triggers that require committee decision before rollback are a profit decline of 2 to 5% for two weeks, a volume collapse exceeding 40%, and a customer complaint rate more than three times the baseline rate.

Metric	**Threshold**	**Action**
Profit decline	>5% for 2 consecutive weeks	Immediate rollback, no committee required
Seller churn	>10% increase vs baseline	Rollback plus seller survey
Review score decline	>0.2 points vs baseline	Rollback plus quality investigation
Competitive undercut	>50% of products	Price war protocol, rollback plus matching

Automatic Rollback Triggers: No committee approval required.

Metric	**Threshold**	**Process**
Profit decline	2-5% for 2 weeks	Investigate before rollback (may be temporary)
Volume collapse	>40% decline	Check elasticity misestimate vs external shock
Customer complaints	>3 times baseline rate	Review complaint themes, assess reputation risk

Manual Review Triggers: Pricing committee decision required before rollback.

7 Discussion & Business Strategy Integration

Sections 4 through 6 presented empirical findings, methodological approaches, and implementation protocols. This section steps back to examine broader strategic implications: how pricing optimization, customer retention, and marketplace dynamics interact to shape competitive positioning, platform economics, and organizational strategy. The discussion integrates technical findings with business strategy, addresses generalizability, and positions this work within the evolving landscape of data-driven marketplace management.

7.1 Strategic Positioning in Two-Sided Marketplaces

7.1.1 Platform Economics and Value Distribution

Two-sided marketplaces face a fundamental tension: value creation must be distributed between sellers and customers to maintain platform viability (Rochet and Tirole, 2003, 2006). The pricing optimization findings reveal how this balance shifts in the Olist context and where platform and seller incentives diverge.

To illustrate the incentive structure, consider a platform commission of approximately 12% on gross merchandise value, a rate consistent with marketplace aggregator models in this segment. Under this structure the platform captures a fixed share of transaction value while sellers retain the remainder net of costs. The tension this creates becomes concrete when profit-optimal pricing is applied: a 20% Electronics price increase with a 32.8% volume decline affects platform revenue through both channels simultaneously:

$$\Delta\text{Platform Revenue} = 0.12 \times \text{Price} \times \text{Quantity} \implies \Delta = (+20\%) \times (-32.8\%) \approx -19\%$$

This calculation reveals a structural misalignment: profit-optimal pricing for sellers reduces platform commission revenue in elastic categories because volume losses more than offset price gains. A platform earning commission on GMV has an implicit preference for volume over margin, which conflicts with seller profit optimization.

This misalignment is real but should not be overstated. Higher seller margins may increase seller retention, reduce platform churn, and generate more sustainable long-run commission revenue than a volume-focused strategy that drives sellers toward unprofitable pricing. The short-run revenue loss from Electronics optimization is approximately 19% of Electronics commission, which represents a modest fraction of total platform revenue, and may be offset by the retention benefits of a healthier seller base.

Three alternative commission structures could better align platform and seller incentives. A profit-sharing model (platform takes a percentage of seller profit rather than GMV) would make the platform a genuine partner in margin optimization, but requires observing seller costs which are currently unobserved. A tiered commission by category elasticity (lower rates on elastic categories, higher on inelastic) would reduce the seller disincentive to raise prices where demand is sensitive. A volume-adjusted commission structure that decreases rates as prices rise within a category would neutralize platform revenue loss from volume declines. Each involves trade-offs between alignment, observability, and administrative complexity. The appropriate near-term recommendation for Olist is to avoid hard pricing guardrails that would undermine seller autonomy, and instead invest in seller-facing analytics tools that make profit-aware pricing the path of least resistance rather than a platform mandate.

Olist's position as a marketplace aggregator distributing seller listings across Mercado Livre, Amazon Brazil, B2W, and Via Varejo simultaneously creates a structurally distinct competitive environment from direct marketplaces. The 96.3% single-seller product rate means most sellers face no direct product-level competition, enabling pricing optimization with limited competitive response risk. This is fundamentally different from Amazon US (where 15 to 20% of products have multiple sellers competing for the buy box) or eBay (where algorithmic repricing creates rapid price convergence). Pricing recommendations derived here may overstate achievable gains in higher-competition marketplace environments where competitor response would dampen the volume-margin trade-off.

7.1.2 Platform Intervention Strategies

Given the incentive misalignment between platform commission revenue and seller profit optimization, platforms have three primary levers for constructive intervention. First, dynamic commission structures tied to category elasticity or seller performance tier would reduce the implicit platform preference for volume. Testing tiered commission rates in the three focus categories (Watches/Gifts, Garden Tools, Electronics) would generate direct evidence on seller price response to commission changes. Second, algorithmic pricing recommendations delivered through seller-facing dashboards provide data-driven guidance without mandating compliance. A pricing optimization service showing sellers their category elasticity, optimal price range, and breakeven cost ratio would operationalize the Section 5 findings as a platform product rather than an academic exercise. Third, A/B test facilitation, where the platform runs pricing experiments on behalf of sellers and shares results, addresses the information asymmetry that prevents individual sellers from estimating their own elasticities given their small transaction samples.

7.2 Competitive Dynamics & Pricing Power

7.2.1 Seller Concentration and Market Structure

The extreme seller concentration documented in Appendix B (Gini coefficient 0.75, top 20% of sellers generating 82% of revenue) has direct implications for how pricing power operates in this marketplace. Single-seller dominance on 96.3% of products means that demand curves facing individual sellers reflect category-level substitution (customers choosing between different products in the same category) rather than seller-level competition (customers choosing between sellers offering identical products). The smooth log-log relationships in the elasticity estimates are consistent with monopolistic competition at the product level rather than the sharp demand responses characteristic of competitive markets.

This concentration enables pricing optimization with limited competitive response risk but creates a strategic paradox for the platform. Seller concentration is simultaneously a source of pricing power (enabling the profit gains documented in Section 5.2) and a platform vulnerability (if top sellers exit, platform revenue collapses disproportionately). The retention focus recommended in Section 5.6 for high-propensity customers applies with equal force to high-revenue sellers: the top 20% of sellers warrant differentiated treatment including tiered commissions, dedicated account management, and early access to pricing analytics tools.

Pricing power in this marketplace is bounded by category elasticity rather than competitor reactions. Electronics buyers compare across products within the category (laptops versus tablets versus smartphones) and exhibit $\eta = -2.18$ price sensitivity regardless of whether any specific competitor undercuts. Garden Tools buyers respond to price with $\eta = -2.77$ reflecting the seasonal and deferrable nature of outdoor purchases. These elasticities represent the structural pricing constraint, and the profit optimization framework operates within them rather than against competitor strategies.

7.3 Customer Lifecycle Economics

7.3.1 Retention Reality and Strategic Targeting

The 3% repeat purchase rate documented in Section 5.5 is low even relative to general merchandise e-commerce benchmarks of 10 to 20%, and dramatically below consumables (40 to 60%) or subscription models (60 to 80%). Three structural factors explain this gap: the marketplace aggregator model prevents customers from forming Olist-specific loyalty since they shop on Mercado Livre or Amazon Brazil without recognizing the underlying platform; the category mix is heavily weighted toward one-time purchase occasions (furniture, electronics, gifts); and the absence of switching costs means customers face zero friction in using alternative platforms for subsequent purchases.

Attempting to close the gap from 3% to consumer-goods benchmarks through retention investment alone is not viable. The structural barriers are too high and the ROI too low for the low-propensity segment, as Section 5.6 demonstrated. The appropriate strategic target is moving high-propensity customers from 4.4% to 6 to 7% repeat rate through loyalty programs, free shipping thresholds, and sequential recommendations, while accepting churn among low-propensity one-time buyers whose first purchases are profitable as standalone transactions.

7.3.2 The CLV Paradox and Retention ROI

The counter-intuitive CLV structure documented in Section 5.6 requires careful interpretation for resource allocation decisions. Low-propensity customers contribute 64% of total CLV through high first-purchase values (BRL 178 average) despite a 2.4% repeat rate. Traditional CLV optimization would direct retention investment toward these high-CLV customers. This is the wrong approach in this setting.

The correct metric is lifetime profitability rather than lifetime revenue. If high-propensity customers (home essentials, multi-item baskets) generate approximately 35% gross margins and low-propensity customers (electronics, furniture, one-time gifts) generate approximately 25% margins, the lifetime profit comparison reverses:

High-propensity lifetime profit: 4.4% $\times$ BRL 106 $\times$ 0.35 $\approx$ BRL 1.63 Low-propensity lifetime profit: 2.4% $\times$ BRL 178 $\times$ 0.25 $\approx$ BRL 1.07

These margin assumptions are illustrative rather than data-derived, but the directional conclusion is robust: high-propensity customers deliver more lifetime profit per acquired customer despite lower CLV, because their purchase occasions recur and their category margins are more favorable. Acquisition strategy should target multi-item home essential buyers over single-item big-ticket buyers, and retention investment should concentrate where repeat probability can realistically be improved.

7.4 Category Portfolio Strategy

The 88% combined category loyalty rate (93% within-session bucket loyalty, 95% within-bucket category loyalty) has a direct and underappreciated strategic implication: category-level pricing optimization does not require portfolio coordination. Because customers so rarely purchase across categories in the same session or across sessions, raising Electronics prices does not cannibalize Garden Tools demand and vice versa. Each category operates as an effectively independent business unit for pricing purposes, which simplifies implementation substantially.

The concentration of profit opportunity in Electronics (82% of total projected gain) combined with this category independence justifies the sequenced rollout recommended in Section 6: validate and optimize Electronics first, since capturing 82% of available profit does not require getting the other categories right. Watches/Gifts and Garden Tools confirmation tests are valuable for completeness but not necessary for the majority of the profit improvement.

Category expansion deserves similar treatment. Given 88% loyalty, adding new categories attracts new customers through acquisition rather than cross-selling existing customers. A new baby and kids category would draw parents entering via search, not Electronics buyers returning for a second category. Category expansion is therefore a customer acquisition strategy, not a retention strategy, and should be evaluated on acquisition cost and new customer lifetime value rather than on cross-sell potential.

The long-tail problem (71 categories, only 4 with robust elasticities) suggests a two-phase approach. In the near term, manual optimization for the top 10 categories by revenue captures the Pareto majority of the opportunity. Over a longer horizon, the methodology developed here could be productized into seller-facing self-service tools that automate elasticity estimation, cost validation prompts, and profit-optimal price recommendations for the remaining categories without requiring analyst intervention on each one.

7.5 Limitations and Generalizability

Four findings from this analysis are likely to generalize beyond this specific context. The aggregation bias result (bucket-level fails, category-level succeeds) reflects a fundamental econometric principle about the unit of analysis that applies wherever demand is estimated from observational data. The cost uncertainty dominance over elasticity uncertainty follows mathematically from the Lerner Index structure and holds whenever margins are thin and demand is elastic. The revenue-profit divergence result is a direct consequence of elastic demand combined with positive marginal costs and applies to any marketplace where sellers face elastic demand at thin margins. The similar magnitude of freight and price effects on retention ($\beta_{\text{freight}} = -0.076, \beta_{\text{price}} = -0.079$) likely generalizes to any marketplace where shipping costs are a material fraction of total purchase cost, which is common in geographically large developing markets.

Four findings are more context-dependent and require local validation before application elsewhere. The 88% category loyalty rate may be specific to Olist's transactional aggregator model and would likely be lower on destination sites or platforms with stronger brand identity. The 3% repeat rate reflects the combination of category mix and aggregator structure and would differ substantially for consumables, subscriptions, or platforms with higher brand recognition. The specific elasticity estimates ($\eta = -2.98$ for Watches, -2.77 for Garden Tools, -2.18 for Electronics) are point estimates from a specific market, time period, and competitive environment and should not be applied to other contexts without local estimation. The geographic freight premium findings (50 to 80% above Southeast baseline in North and Northeast Brazil) reflect Brazilian logistics infrastructure and would not apply directly to markets with different shipping cost structures.

Practitioners adopting this framework should treat the methodology as transferable and the specific estimates as illustrative. The appropriate sequence is to validate context alignment, apply the analytical framework to local data, run pilot experiments before full deployment, and build internal capability for continuous estimation rather than relying on point estimates from a different market.

7.6 Strategic Imperatives

Five strategic imperatives emerge from integrating the empirical findings with the marketplace dynamics analysis.

Cost validation precedes everything else. A 5-point COGS error forfeits 70 to 75% of potential profit gains, making cost surveys the highest-ROI investment in the entire transformation. No pricing experiment should launch before costs are validated to within ± 2.5 percentage points.

Platform and seller incentives are misaligned under fixed GMV commission structures, and this misalignment is structural rather than incidental. Platforms serious about seller profitability should explore tiered commission models and seller-facing analytics tools rather than expecting sellers to optimize independently with limited data.

Retention investment should be allocated by propensity, not by CLV. High-value one-time buyers are profitable as acquisition targets but represent poor retention investments. High-propensity multi-item buyers are the appropriate focus for loyalty programs and sequential recommendation campaigns.

Category independence enables focused optimization. The 88% loyalty rate means Electronics optimization can proceed without coordinating with Watches/Gifts or Garden Tools, and the 82% profit concentration in Electronics means getting that one category right delivers the majority of available value.

The framework transfers but the estimates do not. The methodology demonstrated here (aggregation-aware elasticity estimation, cost sensitivity analysis, propensity-weighted CLV, sequential transition analysis) is designed to be replicable across marketplace contexts. The specific numbers are conditional on this dataset, time period, and market structure and require local validation before implementation elsewhere.

8 Conclusion

This paper developed an integrated pricing transformation framework for two-sided e-commerce marketplaces, combining demand elasticity estimation, profit optimization under cost uncertainty, customer propensity modeling, and sequential purchase pattern analysis. Applied to 110,840 transactions across 71 product categories from the Olist Brazilian marketplace (2016 to 2018), the analysis demonstrates how data-driven pricing strategy can improve marketplace profitability while navigating the structural tensions between platform economics, seller independence, and customer value.

8.1 Summary of Key Findings

The central empirical finding is that revenue maximization and profit maximization are fundamentally misaligned in elastic markets with material costs. Testing a revenue-maximizing 40% price cut generates +160% revenue but -162% profit, converting profitable operations into sustained losses. The profit-optimal strategy, derived from the Lerner Index, recommends raising Electronics prices by 20% (generating +5.6% profit and BRL 3,143 in monthly gains), modest adjustments for Watches/Gifts (-2%) and Garden Tools (+1.5%), and no action where current pricing is already near-optimal. Electronics accounts for 82% of total optimization opportunity despite representing 16% of transaction volume.

A methodological finding with broad implications is that cost uncertainty dominates elasticity uncertainty in determining pricing strategy outcomes. A 5-point COGS error forfeits 70 to 75% of potential profit gains and can reverse strategic direction entirely. A 20% elasticity error preserves profitability across all scenarios. This finding inverts the typical research emphasis: practitioners should invest first in cost validation through seller surveys (2 to 4 weeks, low cost) rather than elasticity refinement through A/B testing (2 to 3 months, foregone revenue), because cost precision matters more than demand precision when margins are thin.

Aggregation level critically determines whether demand elasticity can be identified from observational data. Bucket-level analysis yields a null result ($\beta = +0.09$, $p = 0.648$) due to product mix shifts masking true price effects. Category-level disaggregation recovers negative elasticities for 65% of categories, with four exhibiting robustness across both simple and controlled specifications: Watches and Gifts ($\eta = -2.98$), Garden Tools ($\eta = -2.77$), Electronics ($\eta = -2.18$), and Consoles and Gaming ($\eta = -1.35$).

Same-order bundling is non-viable in this marketplace (0% co-purchase rate across focus categories), but sequential category exploration is substantial: 53% of repeat customers purchase from a different category in their next order, with a median return time of 29 days. This pattern supports time-based recommendation campaigns triggered at days 21, 29, and 56, replacing "frequently bought together" logic with transition-probability-driven sequential suggestions.

The propensity model establishes that freight costs affect repeat purchase probability at a similar magnitude as product price ($\beta_{\text{freight}} = -0.076$, $\beta_{\text{price}} = -0.079$), making geographic freight subsidies as valuable for retention as price reductions. The CLV analysis reveals a counter-intuitive segmentation: low-propensity customers contribute 64% of total CLV through large first purchases but represent poor retention investments because their purchase occasions do not recur and their structural barriers (high freight, low satisfaction) make retention ROI negative. Retention investment should target high-propensity multi-item buyers who deliver higher lifetime profit despite lower CLV.

8.2 Theoretical Contributions

This paper makes three theoretical contributions to the empirical pricing and marketplace economics literature. First, it provides empirical quantification of the relative importance of cost versus elasticity uncertainty in the Lerner pricing framework, a gap in prior work that focuses almost exclusively on demand estimation precision while treating costs as observable. The finding that cost errors dominate elasticity errors has direct implications for research priorities: sophisticated identification strategies for elasticity may be optimizing the secondary parameter while the primary constraint (cost uncertainty) goes unaddressed.

Second, it integrates demand elasticity estimation, customer lifetime value modeling, and sequential purchase pattern analysis into a unified framework applied to a single dataset. Prior e-commerce research typically examines these components in isolation. The demonstration that freight operates through both demand and retention channels simultaneously (freight-volume correlation $r = -0.576$ and retention coefficient $\beta_{\text{freight}} = -0.079$) is an example of the insights available only from integrated analysis.

Third, it documents aggregation bias in a modern e-commerce context, extending Tellis (1988)'s scanner data findings to marketplace settings. The bucket-level null result despite category-level identification success demonstrates that strategic product groupings designed for merchandising purposes are too coarse for demand identification, a methodological caution relevant to any demand study using platform-defined product hierarchies.

8.3 Practical Contributions

The implementation framework in Section 6 provides an operational roadmap that distinguishes this paper from standard applied economics work. The phased approach (cost validation, category prioritization, A/B testing, full rollout) operationalizes the key finding that cost validation must precede elasticity refinement, and the pre-specified decision rules, interim stopping protocols with Bonferroni correction, and rollback thresholds make the experimental design directly actionable rather than illustrative.

The conditional recommendation structure, presenting pricing guidance as decision trees by cost scenario rather than point estimates, acknowledges parameter uncertainty honestly while remaining actionable. The Electronics recommendation is directionally robust across the entire plausible cost range (breakeven COGS 54%, below any realistic industry minimum) while Watches/Gifts and Garden Tools recommendations are explicitly flagged as sensitive to cost validation outcomes. This treatment of uncertainty as a first-class output rather than a limitation to be minimized in a footnote reflects the practitioner reality that decisions must be made under imperfect information.

The propensity-based retention segmentation framework similarly provides actionable structure: invest in high-propensity customers (top 20%), apply standard nurture to medium-propensity, and accept churn among low-propensity customers whose first purchases are profitable as standalone transactions. This segmentation logic, grounded in the empirical finding that lifetime profit is the relevant metric rather than lifetime revenue, challenges the conventional CLV-weighted targeting that would direct resources toward the wrong customers.

The experimental design includes Bonferroni-corrected interim stopping rules at Months 2 and 3, pre-specified decision thresholds (profit lift above 3% at $p < 0.05$ triggers full rollout), and a minimum detectable effect of 4.0% against an expected treatment effect of 5.6%, giving approximately 88% power.

8.4 Limitations

Three limitations bound what can be claimed from this analysis. Elasticity estimates are conditional correlations rather than causal treatment effects. Endogenous pricing (sellers raising prices in response to positive demand shocks) creates upward bias toward zero, meaning robust negative elasticities are conservative lower bounds on true price sensitivity. The absence of valid instruments prevents two-stage least squares correction. Causal validation requires the A/B testing protocols developed in Section 6.

Cost structure is unobserved. All profit optimization findings are conditional on the 65% COGS baseline validated across a 40 to 80% sensitivity range. The sensitivity analysis demonstrates that 5-point cost errors dominate 20% elasticity errors in their impact on recommendations, making cost validation the essential prerequisite for implementation rather than an optional refinement.

External validity is limited to marketplaces with structural characteristics similar to Olist: two-sided aggregator model, seller pricing autonomy, predominantly single-seller products, and geographically large territory with uneven logistics infrastructure. The methodology transfers across contexts; the specific elasticities, loyalty rates, and repeat purchase frequencies are conditional on this dataset and market structure.

8.5 Future Research Directions

Four directions would extend and strengthen this work. Randomized pricing experiments through the A/B testing framework in Section 6 would establish causal elasticity estimates and quantify the endogeneity bias in the observational estimates. Multi-armed bandit designs would enable dynamic learning during experimentation rather than fixed treatment assignments.

Supply-side analysis of seller cost structures and pricing behavior would address the primary data limitation. Understanding whether top sellers have systematically lower costs, whether sellers employ profit-maximizing or heuristic pricing rules, and how sellers respond to platform pricing recommendations would improve both the framework and adoption rate projections.

Dynamic pricing models that adjust recommendations based on inventory levels, seasonal demand variation, and competitive moves would extend the static Lerner framework to a richer optimization problem. The temporal patterns in the sequential purchase analysis provide a starting point for demand state characterization.

Personalized pricing based on customer propensity scores raises empirical and design questions about optimal price dispersion across segments, consumer reactions to price differentiation, and the role of loyalty programs and coupons as obfuscation mechanisms. The propensity model in Section 5.6 provides the segmentation foundation; the pricing optimization layer remains an open problem.

8.6 Closing Reflections

Three overarching lessons emerge from integrating the empirical findings with the implementation and strategy analysis.

Cost validation is the critical path. Academic pricing research prioritizes elasticity estimation sophistication, but the finding that a 5-point COGS error forfeits 70 to 75% of profit gains while a 20% elasticity error preserves profitability across all scenarios redirects implementation priorities fundamentally. Two to four weeks of seller surveys generate more implementation value than two to three months of pricing experiments when costs are unknown.

Revenue and profit maximization diverge in elastic markets, and platform commission structures compound this divergence. The -162% profit impact from revenue-maximizing pricing and the 19% platform revenue decline from profit-optimal pricing illustrate how volume-focused metrics (GMV, transaction count, market share) create misaligned incentives throughout the marketplace ecosystem. Commission redesign or seller-facing analytics tools are necessary to align platform and seller objectives around sustainable profitability.

Customer lifetime value paradoxes require propensity-weighted rather than CLV-weighted targeting. Low-propensity customers contribute 64% of total CLV through large first purchases but deliver negative retention ROI because their structural barriers cannot be overcome through marketing investment. High-propensity customers deliver 52% higher lifetime profit margin despite lower CLV, because repeat probability combined with favorable category margins generates more value than first-purchase revenue alone.

The implementation framework translates these lessons into a phased operational roadmap acknowledging that observational elasticity estimates require experimental validation, unobserved costs demand seller surveys, and organizational adoption requires gradual implementation with clear success metrics and rollback protocols. The total projected annual gain of BRL 56,508 (8.7% above baseline) is conditional on cost validation, A/B test confirmation, and retention program conversion rates consistent with industry benchmarks. The value of this paper lies not in promising specific outcomes but in providing an honest, rigorous path from empirical analysis to strategic execution in a setting that rewards analytical discipline over confident overstatement.

References

Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics*, 37(3):668–691.

Baye, M. R., Morgan, J., and Scholten, P. (2004). Price dispersion in the small and in the large: Evidence from an internet price comparison site. *Journal of Industrial Economics*, 52(4):463–496.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Brynjolfsson, E. and Smith, M. D. (2000). Frictionless commerce? a comparison of internet and conventional retailers. *Management Science*, 46(4):563–585.

Buckinx, W. and Van den Poel, D. (2005). Customer base analysis: partial defection of behaviourally loyal clients in a non-contractual FMCG retail setting. *European Journal of Operational Research*, 164(1):252–268.

Cachon, G. P., Terwiesch, C., and Xu, Y. (2005). Retail assortment planning in the presence of consumer search. *Manufacturing & Service Operations Management*, 7(4):330–346.

Cavallo, A. (2017). Are online and offline prices similar? evidence from large multi-channel retailers. *American Economic Review*, 107(1):283–303.

Chevalier, J. A. and Goolsbee, A. (2003). Measuring prices and price competition online: Amazon.com and barnesandnoble.com. *Quantitative Marketing and Economics*, 1(2):203–222.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–1859.

Einav, L., Kuchler, T., Levin, J., and Sundaresan, N. (2015). Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics*, 7(2):215–247.

Forman, C., Ghose, A., and Goldfarb, A. (2009). Competition between local and electronic markets: How the benefit of buying online depends on where you live. *Management Science*, 55(1):47–57.

Gupta, S. and Lehmann, D. R. (2005). Managing customers as investments: The strategic value of customers in the long run.

Hagiu, A. and Wright, J. (2015). Marketplace or reseller? *Management Science*, 61(1):184–203.

Hortaçsu, A., Martínez-Jerez, F. A., and Douglas, J. (2009). The geography of trade in online transactions: Evidence from ebay and mercadolibre. *American Economic Journal: Microeconomics*, 1(1):53–74.

Lerner, A. P. (1934). The concept of monopoly and the measurement of monopoly power. *Review of Economic Studies*, 1(3):157–175.

Manchanda, P., Ansari, A., and Gupta, S. (1999). The "shopping basket": A model for multicategory purchase incidence decisions. *Marketing Science*, 18(2):95–114.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029.

Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND Journal of Economics*, 37(3):645–667.

Rysman, M. (2009). The economics of two-sided markets. *Journal of Economic Perspectives*, 23(3):125–143.

Tellis, G. J. (1988). The price elasticity of selective demand: A meta-analysis of econometric models of sales. *Journal of Marketing Research*, 25(4):331–341.

Venkatesan, R. and Kumar, V. (2004). A customer lifetime value framework for customer selection and resource allocation strategy. *Journal of Marketing*, 68(4):106–125.

Appendix A: Product Bucket Construction and Category Mappings

This appendix documents the construction of the 10 strategic product buckets used throughout the analysis. Section 2 introduced the bucket structure and summarized bucket-level performance. The purpose here is to provide the complete category-to-bucket mapping and the rationale for each bucket's composition.

Rationale and Assignment Criteria

The 71 granular product categories were aggregated into 10 buckets for three reasons. Many individual categories have sparse monthly observations insufficient for the fixed effects specifications in Section 5.1. Sequential purchase pattern analysis requires sufficient transition volume for stable probability estimates, which aggregation provides. And platform-level strategic analysis operates at the bucket level, not the micro-category level.

Each category was assigned based on purchase motivation similarity (the underlying consumer need driving the purchase), cross-bucket coherence (buckets reflect meaningful budget trade-offs rather than administrative boundaries), and sufficient transaction volume. The 93% within-session bucket loyalty documented in Section 5.4 validates that the 10 buckets correspond to distinct shopping missions: if boundaries were arbitrary, customers would mix across buckets more freely. The dominant sequential transitions documented in Section 5.5 (LEISURE_LIFESTYLE $\leftrightarrow$ HOME_ESSENTIALS, HOME_ESSENTIALS $\leftrightarrow$ PERSONAL_CARE) correspond to recognizable household need progression rather than statistical artifacts, providing further structural validation.

Category-to-Bucket Mappings

Bucket	Categories	N
LEISURE_LIFESTYLE	sports_leisure, toys, watches_gifts, cool_stuff, books_general_interest, musical_instruments, pet_shop, fashion_bags_accessories, books_technical, books_imported, art, party_supplies, flowers, dvds_blu_ray, music, cds_dvds_musicals	16
HOME_ESSENTIALS	bed_bath_table, housewares, furniture_decor, garden_tools, home_appliances, kitchen_dining_laundry_garden_furniture, home_construction, furniture_living_room, furniture_bedroom, air_conditioning, home_confort, home_comfort_2, furniture_mattress_and_upholstery	13
ELECTRONICS_TECH	computers_accessories, telephony, electronics, computers, consoles_games, audio, tablets_printing_image, fixed_telephony, cine_photo	9
AUTO_TOOLS	auto, construction_tools_construction, construction_tools_safety, costruction_tools_garden, construction_tools_lights, costruction_tools_tools, signaling_and_security, agro_industry_and_commerce, industry_commerce_and_business	9
FASHION_APPAREL	fashion_shoes, fashion_male_clothing, fashion_underwear_beach, fashio_female_clothing, fashion_sport, fashion_childrens_clothes, luggage_accessories	7
PERSONAL_CARE	health_beauty, perfumery, baby, diapers_and_hygiene	4
SMALL_APPLIANCES	small_appliances, small_appliances_home_oven_and_coffee, home_appliances_2	3
FOOD_BEVERAGE	food, drinks, food_drink	3
OFFICE_STATIONERY	stationery, office_furniture	2
MISC	market_place, christmas_supplies, la_cuisine, security_and_services, arts_and_craftmanship	5

LEISURE_LIFESTYLE and HOME_ESSENTIALS are the two largest buckets by revenue (29% and 25% respectively) and contain the three focus categories for profit optimization: watches_gifts and garden_tools in LEISURE_LIFESTYLE and HOME_ESSENTIALS respectively, and electronics in ELECTRONICS_TECH. PERSONAL_CARE is maintained as a separate bucket despite containing only four categories because consumable personal care products have recurring purchase characteristics that would confound home improvement demand patterns if merged into HOME_ESSENTIALS. SMALL_APPLIANCES is similarly separated from HOME_ESSENTIALS because its purchase profile (BRL 354 average, 1% of orders, 3% of revenue) reflects infrequent high-value transactions rather than recurring maintenance purchases. MISC categories collectively represent less than 1% of orders and are excluded from all focus analyses in Sections 5.1 through 5.6.

Appendix B: Marketplace Dynamics & Competitive Structure

This appendix presents detailed quantitative analysis of the Olist marketplace's seller concentration, competitive structure, and geographic pricing dynamics. Section 2 provided a summary overview of these patterns. The purpose here is to document the supporting evidence in sufficient detail to allow replication and to provide the analytical foundation for the strategic claims made in Sections 5 through 7.

Seller Concentration Analysis

The analysis covers 2,970 active sellers (those with at least one delivered order during the observation period), generating 97,819 orders and BRL 13,279,836 in total revenue. The mean order count per seller is 32.9, compared to a median of 7.0, a ratio of 4.7 that signals extreme right skew and concentration among top performers.

Revenue concentration substantially exceeds what a uniform distribution would predict. The top 20% of sellers (594 sellers) generate 82.3% of total revenue (BRL 10,929,014), while the bottom 50% (1,485 sellers) contribute only 3.3% (BRL 438,930). The top individual seller generates BRL 226,988 in revenue (1.71% of the total marketplace), and the top 10 sellers combined generate 13.3%. The Gini coefficient estimated from the Lorenz curve is 0.75, which exceeds the 0.65 to 0.70 range typical of large e-commerce platforms such as Amazon US, confirming that this marketplace exhibits winner-take-most dynamics.

The below table profiles the three seller tiers.

Tier	Sellers	Revenue (BRL)	% Total	Avg Rev/ Seller	Avg Orders/ Seller	Avg Rating	Avg Products
Top 20%	594	10,929,014	82.3%	18,399	126.3	4.09	35.8
Middle 30%	891	1,911,893	14.4%	2,146	19.2	4.13	9.4
Bottom 50%	1,485	438,930	3.3%	296	3.8	4.18	2.7

Note: Top-tier sellers average 126 orders and 36 products, consistent with full-time e-commerce operations. Bottom-tier sellers average 3.8 orders, consistent with hobbyists or platform testers.

A counter-intuitive finding emerges from the rating data: top-tier sellers have the lowest average ratings (4.09) while bottom-tier sellers have the highest (4.18), despite the bottom tier generating 270 times less revenue per seller. The correlation between total revenue and average rating is essentially zero ($r = -0.013$). Three explanations are plausible. First, a volume effect: sellers processing thousands of orders accumulate negative reviews through operational failures at scale, while sellers with four orders can maintain perfect ratings through limited exposure. Second, a statistical regression to the mean: small samples allow extreme ratings (all 5-star) that large samples cannot sustain. Third, a customer expectations effect: customers judge established high-volume sellers more critically than small sellers. The business implication is that the marketplace does not penalize low ratings within the acceptable range of 3.5 to 4.5, suggesting customers prioritize price and selection over review score when making purchase decisions.

The distribution of orders per seller follows a power law consistent with other two-sided marketplaces, with an estimated exponent of $\alpha \approx 2.1$, within the typical range of 1.8 to 2.5 for e-commerce platforms.

Multi-Seller Competition Analysis

Of 32,216 products in the marketplace, 96.3% (31,027) are sold by a single seller. The remaining 3.7% (1,189 products) with multiple sellers generate 13.3% of total revenue, a 3.6x revenue

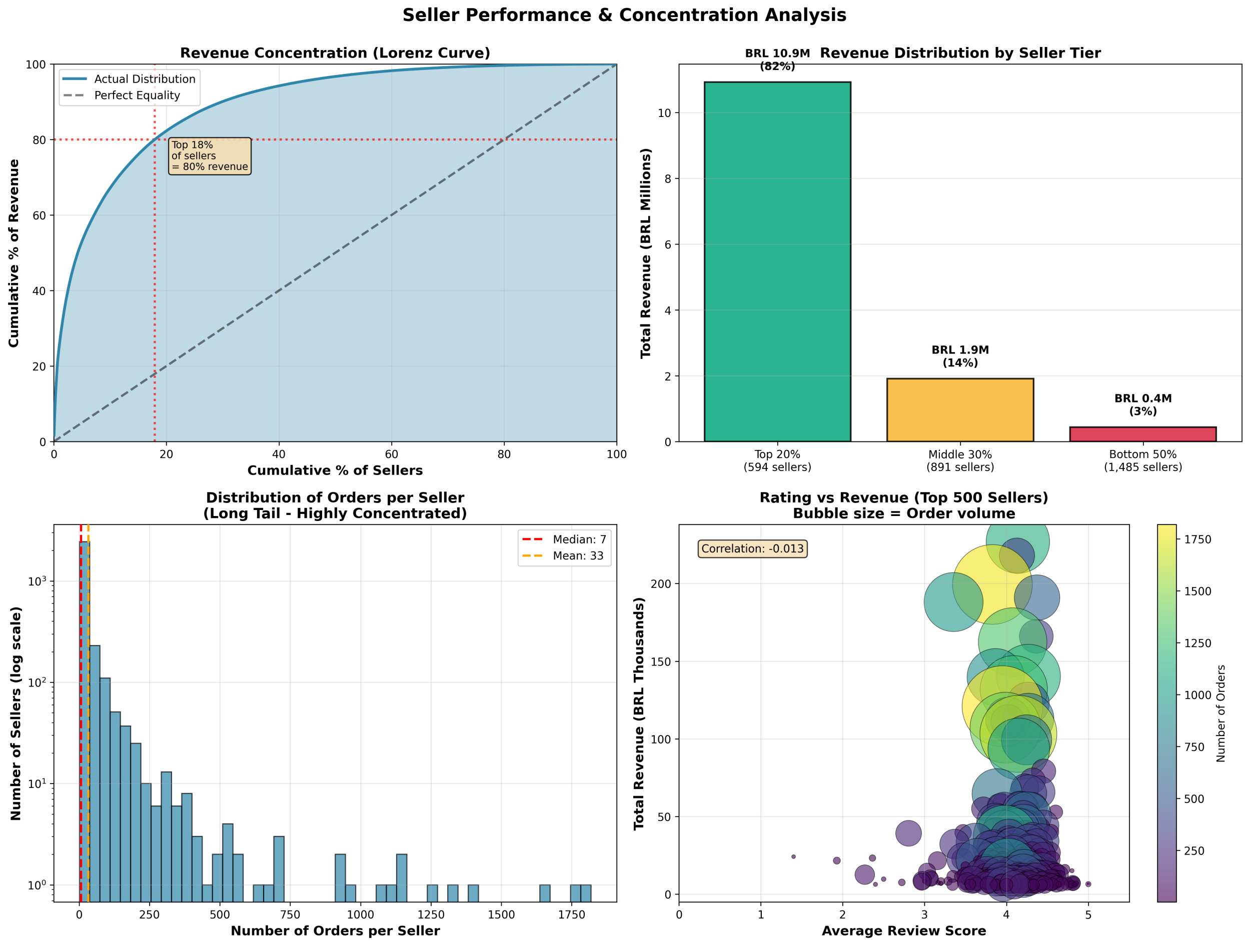

Figure 16: Seller concentration analysis: Lorenz curve (top-left), revenue by tier (top-right), power law order distribution with $\alpha \approx 2.1$ (bottom-left), and inverse rating-revenue relationship with correlation –0.013 (bottom-right).

concentration relative to their product share. This concentration reflects that multi-seller products are high-demand commodities with proven market acceptance.

Among multi-seller products, competition is predominantly duopolistic: 84.0% of competitive products have exactly two sellers, 12.0% have three, and only 4.2% have four or more. Price dispersion is measured using the coefficient of variation (CV = standard deviation / mean price), which normalizes for price level and enables comparison across categories. The overall median CV is 0.061, indicating that most competitive products exhibit prices within 6% of each other.

The distribution of competition intensity is as follows.

Competition Intensity	**Products**	**% of Multi-Seller**
Tight (CV $<$ 0.05)	464	40.3%
Moderate (CV 0.05–0.15)	498	43.3%
Loose (CV 0.15–0.30)	163	14.2%
High dispersion (CV $>$ 0.30)	26	2.3%

Note: 38 products had identical prices across all sellers (CV = 0.000), likely reflecting manufacturer minimum advertised price enforcement or automated repricing to match.

The Duopoly Coordination Effect: A striking finding is that price variance increases rather than decreases as the number of sellers rises, directly contradicting the standard prediction that more competition leads to price convergence.

Number of Sellers	Products	Mean CV	Median CV
2 (duopoly)	1,030	0.083	0.058
3	147	0.098	0.070
4	32	0.098	0.086
5+ (oligopoly)	20	0.112	0.108

Note: Median CV increases 86% from duopoly (0.058) to oligopoly (0.108), contradicting the standard prediction that more competition narrows price variance.

Duopolies coordinate at a median CV of 0.058 while oligopolies fragment at 0.108. Three mechanisms likely explain this pattern. In duopolies, two sellers can monitor each other's prices directly and achieve tacit coordination through mutual interdependence, a Cournot equilibrium outcome. As seller count rises, monitoring costs increase, one aggressive pricer can destabilize the equilibrium, and quality stratification emerges as sellers differentiate across price tiers rather than competing head-to-head. This finding suggests that the platform should encourage three or more sellers per product to break duopoly coordination rather than assuming that more sellers automatically produce more competitive pricing.

Connection to Elasticity Findings: The four categories with robust negative elasticities (Watches/Gifts $\eta = -2.98$, Garden Tools $\eta = -2.77$, Electronics $\eta = -2.18$, Consoles $\eta = -1.35$) exhibit moderate competition (CV 0.73 to 0.86), confirming that elasticity reflects within-category substitution across products rather than seller-level price competition. Categories with positive elasticities (bed_bath_table +8.26, furniture_decor +4.60) exhibit tight competition (CV 0.49), consistent with quality signaling dynamics where tight price bands and brand differentiation allow higher prices to attract rather than repel demand.

Geographic Pricing Dynamics

The marketplace operates as a São Paulo-centric hub-and-spoke network. São Paulo state accounts for 59.6% of all sellers and 64.4% of total revenue, while 63.8% of transactions cross state lines. This structure creates systematic geographic pricing disparities driven primarily by freight costs.

Regional Pricing Summary:

Region	Avg Product Price (BRL)	Avg Freight (BRL)	Total Cost (BRL)	Orders	Customers
North	161.70	36.74	198.45	2,026	1,796
Northeast	148.09	32.23	180.32	10,137	9,044
Center-West	130.50	22.97	153.47	6,541	5,624
South	119.61	21.18	140.79	15,967	13,814
Southeast	114.06	17.33	131.39	76,169	66,200

Note: North customers pay 51% more in total cost than Southeast customers (BRL 198 versus BRL 131). Freight accounts for 13 to 18% of total cost depending on region.

The Pearson correlation between average freight cost and order volume across states is $r = -0.576$ ($p < 0.001$), confirming that freight cost is a primary demand barrier. States with the highest freight costs (Paraíba BRL 43, Rondônia BRL 41, Acre BRL 40) have the lowest order volumes. States with the lowest freight costs (São Paulo BRL 17, Minas Gerais BRL 21, Rio Grande do Sul BRL 21) generate the overwhelming majority of transactions.

Cross-state transactions reveal an important selection effect. Customers purchasing cross-state pay 30% more in total cost (BRL 153 versus BRL 117 for same-state), and notably, the product price itself is 25% higher (BRL 129 versus BRL 104) rather than just the freight component. This reflects selection bias: customers only order cross-state when the product is unavailable locally or represents sufficient value to justify the shipping cost, which skews the cross-state sample toward higher-value purchases.

Connection to Propensity Model: The propensity model in Section 5.6 establishes that freight and price affect retention probability at approximately similar magnitude ($\beta_{\text{freight}} = -0.076, \beta_{\text{price}} = -0.079$). The geographic freight distribution quantifies the retention penalty for remote customers. A North region customer facing BRL 37 average freight versus BRL 17 in the Southeast incurs a freight differential of BRL 20, which translates through the propensity coefficient to approximately a 14% reduction in relative repeat probability. This compounds the acquisition barrier (freight correlation $r = -0.576$ with order volume) to create a dual lifecycle disadvantage for remote customers that standard demand analyses focused on product price alone would miss entirely.

Strategic Implications: Three geographic interventions emerge from this analysis. Free shipping thresholds on orders above BRL 150 in North and Northeast regions would reduce the freight barrier for the majority of qualifying purchases given the higher average product prices in those regions (BRL 148 to BRL 162). Regional pricing tiers offering 15% product discounts in remote regions would approximately equalize total delivered cost between Southeast (BRL 131) and Northeast/North (BRL 138 after discount), bringing these regions into the competitive range without eliminating margin. Local seller recruitment targeting Rio de Janeiro, which exhibits a significant supply-demand gap (13.3% of customer revenue but only 6.2% of sellers), would reduce cross-state freight costs for the São Paulo to Rio corridor which accounts for 9,458 transactions, the largest single cross-state flow.

Integration with Main Paper Findings

The marketplace structure analysis in this appendix supports three claims made in the main text that require the detailed evidence presented here. First, the 96.3% single-seller product rate validates the assumption in Section 5 that elasticity estimates reflect category-level substitution rather than seller-level competition, making the Lerner pricing framework appropriate for this setting. Second, the duopoly coordination finding explains why the competitive response risk assessment in Section

6 treats individual seller price changes as unlikely to trigger systematic competitor reactions. Third, the similar magnitude of freight and price in the propensity model ($\beta = -0.079$ for price and $\beta = -0.076$ for freight) is grounded in the geographic freight distribution documented here, which shows that freight represents 13 to 18% of total customer cost and varies by up to 150% across regions, making it a material decision variable rather than a minor transaction cost.

The combined geographic and propensity analysis also explains the segmentation pattern documented in Section 5.6: high-propensity customers (low freight BRL 17, high satisfaction 4.28 reviews) are structurally concentrated in the Southeast, while low-propensity customers (high freight BRL 24, lower satisfaction 3.92 reviews) are disproportionately from remote regions where freight barriers and selection-biased product mix create structural retention disadvantages that marketing investment cannot overcome.