UNIVERSITAT D'ALACANT
UNIVERSIDAD DE ALICANTE
Facultat de Ciències Econòmiques i Empresarials
Facultad de Ciencias Económicas y Empresariales

MÁSTER EN "ECONOMICS AND DATA SCIENCE"

CURSO ACADÉMICO 2025 - 2026

"The Causal Impact of Logistics Frictions on E-Commerce: Markup Absorption and Channel Substitution"

LUIS MARTÍNEZ VICENTE

TUTORA: ANDREEA CORINA DOBRE (CENTAURO)

CONTENTS

Abstract

This thesis estimates the causal impact of logistics shocks on e-commerce using a high-frequency micro-panel from the Brazilian marketplace Olist. To overcome the extreme persistence of prices and cross-sectional dependence inherent to digital networks, the empirical strategy implements Lag-augmented Local Projections controlled by Two-Way Fixed Effects (TWFE). Findings reveal an absolute nominal rigidity across all segments. In light goods, algorithmic visibility rules prevent cost pass-through, forcing sellers into markup absorption. Conversely, in heavy and bulky goods where margins cannot absorb the shock, elevated freight costs trigger asymmetric demand destruction, driving a permanent offline channel substitution (offline flight).

Keywords: E-commerce, Local Projections, Nominal Rigidity, Logistics

1 INTRODUCTION

1.1 Context and Motivation: Logistics Friction in E-Commerce

Over the past decade, e-commerce has undergone an unprecedented global expansion, reshaping the structure of traditional *retail*. Yet as the market has matured, it has become increasingly clear that the financial viability of digital platforms depends not only on user acquisition but also on the structural efficiency of their supply chains. Logistics has therefore ceased to be merely a cost center and has become a central vulnerability threatening the profitability of the small and medium-sized enterprises (SMEs) operating on centralized *marketplaces*.

This dynamic is particularly acute in emerging economies with infrastructure deficits. According to the Economic Commission for Latin America and the Caribbean (ECLAC), systemic inefficiencies cause logistics costs in the region to absorb between 18% and 32% of a product's total value, a disproportionate burden relative to the 8% average recorded in the United States (InvasWMS, 2024). In Brazil —the geographic setting of this study— national logistics costs amount to 12.5% of Gross Domestic Product (GDP). This burden is compounded by the country's heavy reliance on road freight, which carries 62% of cargo and generates substantial additional energy costs (Cargas e Transportes, 2022).

In *e-commerce*, this macroeconomic inefficiency collides directly with the most complex segment of distribution: *last-mile delivery*. Industry research indicates that, driven by inflation, last-mile delivery now accounts for 53% of total shipping costs, a substantial increase from the 41% recorded in 2018 (Locus, 2024; MegaMinds, 2024).

This structural increase in logistics costs inevitably reaches the consumer interface, giving rise to a behavior known as *Landed Cost Aversion*. Unexpected surcharges during *checkout* are among the most important sources of digital demand destruction. Longitudinal evidence indicates that the global average cart-abandonment rate reaches 70.22%, with "extra costs too high" constituting the leading reason for abandonment in 40% of cases (Baymard Institute, 2024).

Faced with consumer-driven demand hyperelasticity and a logistics cost *shock* imposed by macroeconomic conditions, sellers confront a complex industrial-organization problem. The literature and empirical evidence suggest that, on platforms governed by hierarchical ranking algorithms (such as Amazon's or Olist's *Buy Box*), sellers have little effective freedom to pass these costs through to prices. The risk of losing organic visibility pushes them toward forced markup absorption and nominal rigidity.

1.2 Research Question and Methodology

Against this backdrop of mounting logistics pressure, this Master's thesis asks the following empirical research question: How do sellers in highly competitive markets absorb an exogenous logistics *shock* on a centralized e-commerce platform, and how does freight alter demand as a function of a product's physical characteristics?

To isolate and quantify this transmission mechanism, the study departs from traditional aggregate macroeconomic approaches and adopts a microeconometric causal-inference framework. It constructs a dynamic, high-frequency (daily) longitudinal panel at the Stock Keeping Unit (SKU) level using the *Brazilian E-Commerce Public Dataset by Olist*, which contains tens of thousands of actual transactions.

To mitigate the biases endemic to this type of data —most notably the cross-sectional interdependence generated by logistics *shocks* that affect the entire network simultaneously— the analysis implements the *Local Projections* methodology proposed by Jordà (2005). Estimation is rigorously controlled using *Two-Way Fixed Effects* (TWFE) and robust Newey-West covariances. This framework isolates the exogenous effect of the logistics *shock* on two response variables: the seller's posted price and final conversion (demand). In addition, to capture treatment heterogeneity, the analysis is divided into two taxonomic clusters defined by their physical cost structure: light goods (technology and impulse-purchase items) and heavy goods (furniture and bulky products).

1.3 Main Findings

The empirical analysis yields two central findings that reveal important frictions in digital markets:

- **Nominal Rigidity Across Segments (*Markup Absorption*):** The results show an almost complete absence of *cost pass-through* into prices. Regardless of whether a product belongs to the light-goods cluster (high margin) or the heavy-goods cluster (low margin), the effect of higher logistics costs on the base price is not statistically significant over a projected 15-day horizon. Empirically, platform-ranking pressure forces sellers to absorb the shock, sacrificing operating margins in order to preserve commercial viability.

- **Asymmetric Demand Destruction (*Channel Substitution*):** Although sellers absorb the cost in order to protect the product's posted price, freight salience at *checkout* substantially erodes conversion. For light goods, the decline in sales is transitory and largely dissipates by around the ninth day; for heavy goods, the effect is permanent. The surcharge induced by dimensional-weight pricing crosses consumers' tolerance threshold, generating a lasting "scarring effect" that encourages substitution from the *online* channel toward traditional physical *retail*.

1.4 Structure of the Paper

The remainder of the paper is organized as follows. Section 2 reviews the recent literature on spatial frictions, algorithmic pricing, and Local Projections. Section 3 describes the Olist dataset, the cluster definitions, and the reengineering of the longitudinal panel. Section 4 presents the econometric strategy and motivates pragmatic modeling choices, including the use of the inverse hyperbolic sine (IHS) transformation. Section 5 subjects the panel to diagnostic tests (Pesaran, Hausman, and unit-root tests). Section 6 presents the univariate empirical results and discusses treatment heterogeneity. Finally, Section 7 synthesizes the economic conclusions and evaluates the corporate implications for the future viability of *e-commerce*.

2 LITERATURE

Analyzing the transmission of logistics *shocks* in digital ecosystems requires insights from several branches of economics. This review organizes the theoretical framework around three core analytical dimensions: (1) economic geography and consumer responses to transportation frictions; (2) contemporary industrial organization and the effects of prominence-allocation algorithms; and (3) recent developments in time-series econometrics for persistent variables in dynamic models.

2.1 Spatial Frictions and Channel Substitution (*Offline Flight*)

The literature on the microstructure of digital markets has extensively examined how consumers process transaction costs in e-commerce relative to traditional retail. Unlike in physical retail, where transportation costs are internalized as the opportunity cost of travel time and distance, in *e-commerce* this spatial friction appears explicitly and prominently as a *shipping fee*. The theory of *partitioned pricing* suggests that consumers evaluate the product's base price and add-on charges at the end of the conversion funnel asymmetrically. Because shipping fees add no intrinsic value to the core product, they generate a salient disutility commonly described as *Landed Cost Aversion*.

This aversion is stronger for bulky or heavy goods. Logistics operators rely on *dimensional weight pricing*, causing supply-chain inefficiencies to affect products with high physical frictions disproportionately. When consumers reach *checkout*, the sudden *salience* of this cost acts as a severe negative *shock* to expected utility, triggering a phenomenon in economic geography known as *offline flight*, or channel substitution.

Recent work by Kanuri et al. (2025) directly examines this cross-channel behavior. Using a Regression Discontinuity in Time (RDiT) design on a panel of more than 21,000 buyers, the authors show that policies that make friction-intensive orders more expensive divert consumers away from the digital ecosystem, documenting a 23% increase in physical-retail sales. They identify an *offline store interpurchase effect*, whereby aversion to shipping fees induces a direct reallocation of spending toward *brick-and-mortar* outlets.

Complementing this evidence from an operations perspective, Wang et al. (2023) develop a game-theoretic model to assess how shipping fees shape the competitive equilibrium between channel integration (omnichannel models such as *Buy Online, Pick Up In Store*, or BOPS) and complete channel separation. Their analytical results show that the shipping fee serves as a signaling mechanism for spatial inefficiency. Because platforms such as Olist operate strictly as pure digital intermediary *marketplaces*, where third-party sellers (*3P sellers*) do not provide omnichannel infrastructure, the ecosystem operates under full separation. In the absence of in-store pickup, consumer disutility is maximized, providing a theoretical basis for why logistics *shocks* in heavy goods translate into a sharp contraction of digital-channel demand.

2.2 Algorithmic Governance, *Buy Box*, and Nominal Rigidity

The second pillar of this study lies in contemporary Industrial Organization and focuses on how digitalization has altered the mechanism of *cost pass-through*. Neoclassical microeconomics typically assumes that an exogenous *shock* that raises marginal costs will be passed through, at least partially, to the price paid by the final consumer. Modern *marketplaces*, however, operate under an algorithmic-governance architecture in which tools that allocate visual prominence —such as the *Buy Box* or *Featured Offer*— act as key *gatekeepers* of demand. Empirical evidence suggests that the overwhelming majority of transactions in these ecosystems are completed through the seller holding this position.

The influential study by Brown and MacKay (2023) reshapes the understanding of platform competition by documenting how algorithmic pricing technology moves the market away from the standard Nash-Bertrand benchmark. The authors show that pricing algorithms function as powerful devices of *short-run commitment*, forcing firms to adopt reactive strategies tightly constrained by platform rules. This dynamic, referred to as algorithmic coercion, creates substantial frictions in the adjustment of retail prices to fluctuations in the underlying cost structure.

To characterize this architecture more precisely, Thakurele et al. (2024) formalize the logic of the *Buy Box* and restrictive price-parity policies. They examine how eligibility for prominence is inseparably linked to an *eligibility-linked price parity rule*. Within their framework, this rule forces sellers to face highly discontinuous profit functions. If an independent seller passes through a cost *shock* by raising its posted price, the algorithm automatically penalizes the offer through *price suppression*. Thakurele et al. (2024) show that this high *compliance cost* pushes sellers toward a suboptimal equilibrium of *markup absorption*: merchants reduce unit profits rather than risk commercial invisibility, thereby inducing pronounced nominal rigidity from the consumer's perspective.

2.3 Inference in Local Projections with Persistent Variables

The third pillar of this review moves beyond economic theory to the frontier of econometric methodology and time-series analysis. Historically, Impulse Response Functions (IRFs) to structural *shocks* have been estimated primarily through Vector Autoregressions (VARs). The Local Projections (LP) method introduced by Jordà (2005), however, has gained broad prominence because of its robustness to misspecification.

The main econometric challenge posed by high-frequency *e-commerce* time series (such as daily prices or sparse sales) is their extreme persistence, often characterized by order-one integration $I(1)$ or *near-unit roots*. In classical econometrics, the conventional response is to take first differences to induce $I(0)$ stationarity. This approach, however, entails important empirical costs: it removes long-run cointegrating relationships and obscures the direct economic interpretation of elasticities estimated in levels.

The paper by Montiel Olea and Plagborg-Møller (2021) provides a fundamental methodological framework for resolving this trade-off. The authors show analytically that Local Projection inference can robustly accommodate extreme persistence in the data-generating process. Their key contribution is the formal development of *Lag-augmented Local Projections*. They establish asymptotically that, when a standard local projection includes additional lags of the system variables as controls, the resulting statistics deliver *uniformly valid inference* regardless of the underlying degree of persistence.

The econometric mechanism underlying this result relies on the Frisch-Waugh-Lovell (FWL) theorem and orthogonal projection. Consider a highly persistent autoregressive data-generating process with a unit root (where $\rho \approx 1$ or $\rho = 1$):

$$y_t = \rho y_{t-1} + u_t \tag{1}$$

If a local projection for horizon h is specified as:

$$y_{t+h} = \beta_h x_t + \Gamma_h' z_t + \epsilon_{t,h} \tag{2}$$

and the control vector z_t is empty, Ordinary Least Squares (OLS) inherits the nonstandard properties of a random walk. By contrast, proactively including lags (such as y_{t-1}) in z_t causes the orthogonal projection of the regressor onto its own lags to isolate the structural innovation component. The effective residualized regressor takes the form:

$$\tilde{y}_t \approx y_t - \rho y_{t-1} = u_t \tag{3}$$

Because the stochastic innovation u_t is a stationary $I(0)$ process, the empirical estimator converges asymptotically to a standard Normal distribution throughout the parameter space.

This result is especially important for applied econometrics for two reasons. First, it avoids overdifferencing and provides the mathematical basis for estimating models directly in levels even when the data are $I(1)$, thereby preserving interpretability. Second, it simplifies the use of HAC/HAR corrections. The authors show that adding lags reduces serial correlation in the regression's asymptotic *scores*, making elaborate heteroskedasticity-and-autocorrelation-consistent (HAC) corrections less essential for addressing overlapping projection horizons.

This theorem provides the econometric foundation for the strategy implemented in this thesis. The large Olist panel contains highly persistent variables: nominal prices display an inertial stepwise pattern consistent with algorithmic rigidity, while daily sales series exhibit strongly $I(1)$ dynamics.

Had this study relied on a conventional differenced VAR, the mathematical transformation of prices would have discarded relevant information about the absolute price level paid by consumers, a key variable for capturing landed cost aversion. Differencing would also have artificially inflated the error variance. Following Montiel Olea and Plagborg-Møller (2021), the Local Projections in this study are estimated in nominal levels and include lags of logistics *shocks* and historical endogenous variables in the vector of exogenous controls. This *Lag-augmentation* acts as a filter that isolates the exogenous logistics *shock*, ensuring asymptotically correct confidence-interval coverage. As a result, the estimates of markup absorption and demand contraction achieve a high degree of inferential robustness.

3 DATA AND PANEL REENGINEERING

The econometric analysis of logistics *shocks* in e-commerce requires a data architecture capable of capturing spatial frictions at the microeconomic level. Transforming a repository of raw transaction data into a longitudinal panel suitable for causal inference is not merely an exercise in data manipulation; it requires a set of fundamental theoretical assumptions that ensure the consistency and unbiasedness of the subsequent estimates.

This section describes the primary data source, motivates the clustering procedure, and justifies the four econometric reengineering (ETL) guidelines designed to mitigate structural biases.

3.1 Presentation and Nature of the Data: The Olist Ecosystem

The empirical analysis is based on the *Brazilian E-Commerce Public Dataset by Olist*, a large anonymized transaction-level database documenting more than 100,000 orders placed in Brazil between 2016 and 2018. Like platforms such as Amazon, Olist operates as a centralized ecosystem whose structural role is to allow small and medium-sized enterprises (SMEs) to integrate and manage their product catalogs across the country's major *marketplaces*. Data from a multicategory platform of this scale provide unusually rich transactional heterogeneity for applied economics, allowing the simultaneous observation of the microeconomic behavior of thousands of sellers and buyers.

The Brazilian market, together with the penetration of this type of *e-commerce* platform, provides an especially useful natural laboratory for evaluating the *shocks* studied here. The rationale is that the broader logistics environment is characterized by extreme spatial frictions arising from continental distances, substantial asymmetries in public-infrastructure provision, and heavy reliance on road freight. These features expose the distribution network to high volatility in freight costs and to recurrent large exogenous disturbances, including nationwide trucker strikes and abrupt inflationary movements in fuel prices.

In this noisy and volatile environment, traditional macroeconomic approaches that aggregate daily transactions into time-series aggregates face severe methodological limitations because cross-sectional aggregation destroys the granular variation in market microstructure and distorts identification of price-transmission channels. The first challenge, therefore, was to construct a matrix that preserved the integrity of the microdata.

To build this matrix, the processing pipeline performed a strict relational *inner join* across three primary tables in the database:

- **Order Data:** This table provides the temporal backbone of the analysis. A strict filter was deliberately applied to retain only orders with confirmed *delivered* status. From a demand-theory perspective, this restriction is necessary to ensure that the resulting measures reflect actual, realized demand, excluding cancellations, fraud, and returns that would contaminate the estimates. The exact purchase timestamp was also extracted and collapsed to daily frequency to define the time dimension (T) of the panel.

- **Transaction Data:** This intersection yields the microeconomic variables entering the demand equations. Specifically, the analysis isolates each product's unique identifier (SKU), the base price set unilaterally by the seller (P_t), and, critically for this thesis, the exact logistics shipping cost assessed and paid by the consumer at *checkout* (C_t). The latter provides the direct empirical measure of the spatial *shock* borne by the buyer.

- **Physical Product Data:** Unlike conventional commercial analyses, this study extracts the intrinsic physical characteristics of the goods. Nominal weight in grams was collected, and the raw spatial dimensions were combined to obtain each item's exact three-dimensional volume in cubic centimeters, defined geometrically as:

  $$\text{Volume} = \text{Length} \times \text{Width} \times \text{Height} \tag{4}$$

  This geometric and gravimetric information is fundamental for correctly segmenting treatment assignment, as motivated and developed in greater detail in the following subsection.

This initial procedure ensures that the transaction repository acquires a coherent tabular structure around verified dates and observable physical goods, thereby establishing the basis for the subsequent bias-correction procedures.

3.2 Taxonomic Criteria and Clustering by Physical Properties: The Dimensional Weight Pricing Paradigm

Controlling for heterogeneous unobserved effects and constructing analytical clusters is essential when evaluating *shocks* involving instrumental variables or spatial differences. A common methodological practice in *e-commerce* studies is to pool observations cross-sectionally using the platform's predefined categories, collapsing all SKUs into generic commercial labels (e.g., "Watches" or "Garden Furniture").

When the central object of causal inference is the effect of logistics disturbances, however, commercial categorization has no underlying causal relevance. Last-mile logistics infrastructure and transportation-pricing algorithms do not classify freight according to its social use; they respond strictly to geometric and gravimetric constraints.

The global industry standard adopted by logistics integrators and postal services (including Correios in Brazil, Olist's core logistics provider) is *Dimensional Weight Pricing* (DWP). This algorithmic pricing model penalizes occupied space according to the following spatial-pricing function[1]:

$$\text{Volumetric Weight} = \frac{\text{Length} \times \text{Width} \times \text{Height}}{\text{DIM Factor}} \tag{5}$$

At billing, the transportation network compares actual weight (*dead weight*) with dimensional weight and invariably bases the freight charge on the larger of the two measures (*billable weight*).

The contemporary supply-chain management literature supports this institutional feature. Singh and Ardjmand (2021) document in detail how the shift toward dimensional-weight pricing transformed the cost structure of e-commerce. They argue that carriers impose DWP to maximize fleet profitability, making freight costs direct and nonlinear functions of package size. In addition, Jararweh et al. (2025) examine how shipping firms use dimensional weight to address the negative externality created by large but lightweight parcels, showing that a product's three-dimensional structure is a dominant determinant of its exposure to transportation costs.

Grouping products by *marketing* category would mix lightweight but highly bulky goods with dense and compact goods. Following a sharp increase in the cost per cubic meter caused by a capacity shock, bulky goods will absorb a large logistics *shock*, whereas compact goods will be affected only marginally. Pooling this variation under a single commercial label would introduce endogeneity into the residuals of the econometric model and bias the estimator.

Accordingly, this Master's thesis implements a clustering strategy based strictly on numerical measures of physical weight and volume, dividing the analysis into two cohorts that face distinct logistics cost functions:

- **Cluster 1 (Light Goods):** This cluster contains high-density, low-volume categories (5.279 cm^3 and approximately 750 g on average). It includes perfumes,

[1] The DIM Factor (volumetric divisor) is an exogenous constant standardized by the logistics operator (typically set at approximately 6,000 cm^3/kg for standard postal services). This constant acts as the density threshold required by the transportation network to ensure profitable use of fleet capacity, allowing package volume to be converted directly into an equivalent billable weight.

watches, computing products, and beauty goods. Logistics for these products are highly optimized and operate in an environment in which three-dimensional constraints are limited; freight therefore represents a relatively minor friction (approximately 14% of the total price).

- **Cluster 2 (Heavy and Bulky Goods):** This cluster includes household products, furniture, garden tools, and appliances, with average volume and mass well above three times those of Cluster 1 (17.878 cm^3 and 2.418 g). These goods cross the threshold into consolidated freight (LTL - *Less-Than-Truckload*) and therefore face substantial dimensional surcharges. In this cohort, freight imposes a significant burden, accounting for between 20% and 24% of transaction value.

Clustering by physical density isolates groups of products that share a common exposure to macroeconomic fluctuations in national logistics rates. This empirical segmentation produces two independent and causally homogeneous analytical datasets for the panel regressions, supporting unbiased parametric estimation.

3.3 Transition to the Micro-Panel and Unit Value Bias Mitigation

In the preliminary analytical iterations, the study considered estimating the logistics effect using a macroeconomic Structural Vector Autoregression (SVAR), an approach that requires collapsing transactions into broad aggregates. This approach was ultimately rejected after it was found to induce a structural bias in the measurement of retail demand: *Unit Value Bias*.

In an ecosystem such as Olist, where product differentiation is pervasive, using "unit value" (total revenue divided by total quantities within a category) as a *proxy* for price collapses together products that may be substitutes but differ substantially in quality and characteristics. When freight becomes more expensive, rational consumers are well known to alter basket composition, substituting lower-quality alternatives for *premium* goods to offset the surcharge (the Slutsky substitution effect). When this reallocation occurs, a researcher observing only the unit-value index will register a decline driven by a smaller share of premium products and a larger share of inexpensive alternatives. The researcher may therefore incorrectly infer that sellers have reduced prices even if individual prices remained unchanged and only the composition of the consumption basket shifted.

The contemporary literature supports the decision to avoid aggregation. Diewert and Fox (2022) provide a formal framework showing that aggregating sales over time windows causes the resulting index to capture consumer substitution rather than genuine strategic supply-side price adjustments. Likewise, Cavallo (2018) shows that, at a granular level, price adjustment is highly asynchronous; aggregation artificially smooths this variation, creating the appearance of flexible pricing and masking true rigidity[2].

The resulting methodological safeguard is a strict longitudinal panel at the individual Stock Keeping Unit (SKU) level. Tracking each product's unique identifier over time controls for the good's intrinsic characteristics and ensures that any variation in the micro-panel reflects genuine price movements.

3.4 Lifecycle Attrition and the Ghost Zeros Problem

Once the SKU-level panel has been constructed, a new challenge arises from the extreme volatility of product life cycles on *marketplaces*. Over a multiyear window, seller turnover is substantial: products enter the platform, sell, experience *stock* outages, or are discontinued. A standard methodological instinct when estimating Fixed Effects models of this kind is to force a balanced panel by filling temporal gaps with zero sales.

The meaning of an empirical zero, however, must be treated with care. An empirical zero is informative: it represents an observed outcome in which the consumer chose not to purchase the item. By contrast, a zero inserted algorithmically in a period when the product did not yet exist on the platform is a *structural zero*, or *ghost zero*. In such periods, the mathematical probability of purchase was zero regardless of the logistics environment.

Introducing these zeros creates substantial *attenuation bias*: the model interprets demand as remaining at zero even when logistics costs may have fallen, flattening the estimated elasticity coefficients toward zero. Aguiar and Kashaev (2024) provide a recent formalization showing how including products outside the consumer's *choice set* distorts demand estimates. In the Olist setting, forcing a strictly balanced panel would therefore create a matrix dominated by structural zeros and generate a censored-demand problem.

To avoid this bias, the analysis implements *Lifecycle Bounding*, treating the platform's *Choice Set* as dynamic and using the observable interval between a product's first and last sale as the best available estimate of its active exposure window.

Trimming observations outside the demonstrable active life of each product makes it plausible that the retained zero-sales days reflect genuine consumer non-purchases, thereby correctly capturing the marginal propensity to consume in response to system-wide logistics *shocks*.

[2] To illustrate this statistical artifact, suppose several competitors keep their prices rigidly unchanged but implement discrete increases at different points within a period (asynchronous adjustment). A daily aggregate average then generates a smooth upward-sloping curve. The aggregation would incorrectly suggest that the market adjusts prices marginally day by day, concealing the underlying *step-function* nature of price setting at the microeconomic level.

3.5 Network Shock Imputation via Cross-Sectional Means and Common Correlated Effects (CCE)

Despite this temporal bounding of the panel, there remain days with no sales within a product's active life. From the demand perspective (Q_t), these are genuine empirical zeros; for the base price (P_t), carrying forward the last nominal value is acceptable because of short-run price rigidities. The treatment of missing values in the logistics freight variable (C_t), however, poses a critical analytical challenge.

Standard imputation routines (such as last observation carried forward or linear interpolation) artificially smooth exogenous volatility. If a logistics *shock* occurs (for example, a trucker strike), freight costs can spike sharply. Logistics cost is not a static SKU characteristic; it is the manifestation of unobserved common factors that permeate the entire network cross-sectionally at time t.

To address this problem without eliminating exogenous variation in the series, the data design draws on the logic of Common Correlated Effects (CCE) estimators proposed by Pesaran (2006) and extended by Chudik and Pesaran (2015). In this framework, the freight cost of product i at time t is modeled using a multifactor error structure:

$$C_{i,t} = \alpha_i + \lambda_i' f_t + v_{i,t} \tag{6}$$

Where:

- α_i is the product's deterministic fixed effect (e.g., its physical volume or brand).
- f_t is a vector of unobserved common factors (the systemic logistics *shock* on day t).
- λ_i is the product's heterogeneous sensitivity to that *shock*.
- $v_{i,t}$ is the purely idiosyncratic random error term.

The empirical objective is to recover the systemic *shock* signal f_t on days when product i has no transactions. Pesaran (2006) shows asymptotically that, as the cross-sectional dimension of the panel grows ($N \to \infty$), cross-sectional averages of observed variables converge to the unobserved common factor. Specifically, consider the daily cross-sectional mean freight cost across all active products:

$$\bar{C}_t = \frac{1}{N}\sum_{i=1}^{N} C_{i,t} = \bar{\alpha} + \bar{\lambda}' f_t + \bar{v}_t \tag{7}$$

Because the idiosyncratic error vanishes by the Law of Large Numbers ($\bar{v}_t \xrightarrow{p} 0$), the daily cross-sectional mean $\bar{C}_t$ becomes a consistent and asymptotically valid *proxy* for the network-wide logistics *shock* at that point in time (f_t)[3].

In the algorithmic implementation (Python), this cross-sectional mean was computed after excluding missing values and grouping by date (`groupby('date').mean()`). The resulting vector $\bar{C}_t$ was then projected (imputed) onto each individual's missing transaction dates.

To mitigate the multicollinearity concerns discussed by Juodis et al. (2022), the daily cross-sectional mean was computed only within each physical cluster (light goods versus heavy goods). This partition allows for between-cluster heterogeneity and within-cluster homogeneity ($\lambda_i \approx \bar{\lambda}$ within a given cluster). Restricting imputation to physically homogeneous clusters minimizes *level shift bias*, ensuring that the baseline levels of imputed products do not differ sharply from their cohort mean $\bar{\alpha}$. The procedure therefore acts as a regularization mechanism, leaving the resulting panel with a clean exogenous signal free from interpolation bias.

Together, these four components produce a highly robust longitudinal data model and provide the subsequent inferential regressions with a clean matrix and a well-defined exogenous signal.

4 EMPIRICAL METHOD

The empirical evaluation of exogenous interventions in digital markets, and in particular the estimation of the causal effect of logistics *shocks* on pricing policy and demand, requires a high-precision econometric framework. Following the data reengineering described in the previous section, the analysis uses a high-frequency longitudinal micro-panel organized at the Stock Keeping Unit (SKU) level with daily time resolution.

In applied econometrics, a two-dimensional panel —where N denotes the number of individual products and T the number of observed days— offers a substantial analytical advantage over simple time series or cross sections. This structure allows the researcher to observe how the same unit (the same product) responds over time to changes in its environment, facilitating control for unobserved heterogeneity. In particular, it isolates intrinsic product characteristics that do not vary over time (such as brand, physical size, or material quality) and that would otherwise bias the estimates.

The objective of this section, however, is not to establish simple statistical correlations but to identify causal effects. In e-commerce, observing that shipping costs rise while

[3]It is important to note that the structural constant $\bar{\alpha}$ (the mean fixed effect across all products in the cluster) is asymptotically irrelevant for the subsequent dynamic inference. This constant is removed mechanically by the *Within* transformation when individual Fixed Effects intercepts (α_i) are included in the Local Projections regression, allowing the model to isolate the time variation induced solely by the *shock* (f_t).

sales fall is not sufficient to establish a direct causal relationship. Omitted variables or simultaneous dynamics (endogeneity) may account for the same pattern. Moving from descriptive statistics to causal microeconometrics with high-frequency data therefore raises important methodological challenges.

As emphasized in the literature, research designs based on daily-frequency micro-panels simultaneously face three structural challenges that invalidate the simplest canonical approaches:

- **Dynamic model specification under high volatility:** Online prices and demand change erratically. Traditional models that impose a rigid linear structure over time risk propagating estimation errors when the initial specification is even slightly incorrect. The appropriate framework must allow the effect of a *shock* to be projected forward horizon by horizon in a flexible and agnostic manner.
- **Robust statistical inference under Cross-Sectional Dependence:** In an interconnected market such as Brazil, a macroeconomic *shock* (for example, a change in Correios' tariff policy or a nationwide road-freight strike) affects multiple sellers in the Olist network simultaneously. This violates the classical assumption of independence across observations, generating what panel econometrics refers to as *Cross-Sectional Dependence* (CSD). The empirical method must remove this systemic component so that spurious variation is not attributed to the model.
- **Parametric modeling of censored variables:** Microeconomic demand is not continuous. On many days, a given product records exactly zero sales. The mathematics underlying elasticity estimation —which traditionally relies on logarithmic transformations— breaks down in the presence of exact zeros, requiring alternative functional transformations while acknowledging the limitations imposed by their asymptotic properties.

This chapter provides a comprehensive theoretical treatment of these three challenges. Its purpose is to establish the analytical framework needed to justify the modeling choices adopted below.

4.1 Local Projections (LP) Model Specification

The central objective of this study is not to measure the static, contemporaneous effect of a logistics cost increase, but to trace the evolution of its consequences over time. In applied economics, when an exogenous *shock* hits a system (for example, when a product's freight cost suddenly increases on day t), its effects on demand and prices do not disappear immediately. Instead, they propagate over subsequent days and generate a dynamic path. Econometrics summarizes this path through Impulse Response Functions (IRFs).

Historically, IRFs in panel settings have been estimated primarily using Panel Vector Autoregressions (*Panel-VAR*). A VAR assumes a first-order linear process in which the state of the system at $t+1$ depends on its state at t:

$$Y_{t+1} = \Phi Y_t + u_{t+1} \tag{8}$$

To obtain the dynamic response at a distant horizon h, the VAR does not estimate that horizon directly. Instead, it iterates the equation forward recursively, yielding:

$$Y_{t+h} = \Phi^h Y_t + \sum_{j=1}^{h} \Phi^{h-j} u_{t+j} \tag{9}$$

This iterative approach has an important weakness when applied to high-frequency e-commerce microdata. If the initial coefficient matrix Φ is even slightly misspecified (*mis-specification*) —for example, because it imposes linearity when the pricing algorithm is asymmetric—, iteration raises the parameter to the power h (Φ^h), geometrically propagating the error and accumulating substantial bias as the projection horizon increases. In this sense, the VAR effectively assumes that the Data-Generating Process (DGP) is known correctly, an assumption that is difficult to sustain in the volatile *e-commerce* environment.

To avoid this problem, the analysis adopts the *Local Projections* (LP) methodology originally developed by Jordà (2005) and established in the recent literature as a leading tool for causal microeconometrics Jordà (2023). Unlike a VAR, Local Projections do not iterate a single equation forward. Instead, the estimator projects the endogenous variable at $t+h$ directly on the structural *shock* at time t.

The baseline panel specification, estimated separately for the two outcomes of interest (Price and Demand), is:

$$y_{i,t+h} = \alpha_i + \gamma_t + \beta_h C_{i,t} + \sum_{j=1}^{p} \delta_{j,h} y_{i,t-j} + \sum_{j=1}^{p} \theta_{j,h} C_{i,t-j} + \epsilon_{i,t+h} \quad \text{para } h = 0, 1, \ldots, 14 \tag{10}$$

To understand the mechanics of the model and its mapping to transaction-level behavior, it is useful to unpack the components of this parametric equation:

- $y_{i,t+h}$ **(Endogenous Variable):** This term denotes the response of individual product (SKU) i, evaluated h days after the *shock* occurs. In the two separate models, the outcome is either the seller's posted Price or Quantity Demanded. As discussed in Subsection 4.3, the structure of the microdata requires specific functional transformations of y to make estimation feasible.

- $\mathrm{C}_{i,t}$ **(Treatment Variable):** This is the logistics cost (freight) observed for product i at the baseline time t. It is the exogenous disturbance whose effect the analysis seeks to isolate.

- β_h **(Parameter of Interest):** This parameter is the core of the inference. β_h captures the direct Average Treatment Effect (ATE) of the *shock* at the exact horizon h. Estimating Equation (1) independently fifteen times (from $h = 0$ through $h = 14$) produces a sequence of fifteen β coefficients. Plotting this sequence yields the Impulse Response Function (IRF).

- **Historical Controls (*Lag-augmentation*):** LPs allow past control variables (*lags*) to be included flexibly. The specification includes a vector with $p = 2$ temporal *lags* of both the endogenous variable ($y_{i,t-j}$) and historical costs ($C_{i,t-j}$). This inclusion is substantive rather than cosmetic: as discussed below, adding lags in this framework allows highly persistent variables to be estimated without differencing.

- α_i **and** γ_t **(Two-Way Fixed Effects):** These parameters are product-specific (α) and day-specific (γ) intercepts. Their inclusion provides the model's primary protection against omitted-variable bias and network interdependence, as developed in Subsection 4.2.

- $\epsilon_{i,t+h}$ **(Error Term):** This term captures all variation not explained by the model. Under Local Projections, the error projected h days ahead mechanically exhibits serial correlation of order MA($h - 1$), requiring robust standard errors for valid confidence intervals.

The choice of this architecture is grounded in the bias-variance *trade-off* documented in the recent econometric literature. Li et al. (2024) evaluate both methods across empirically calibrated Data-Generating Processes (DGPs). Their results show that although VARs have lower statistical variance at short horizons because of their stronger parametric structure, they can exhibit substantial bias when variables have long memory, near unit roots, or complex dynamics.

Because the daily panel contains erratic fluctuations that are unlikely to admit a perfectly specified VAR representation, Local Projections provide the more suitable specification. By avoiding an iterative global law of motion, the LP estimator remains largely agnostic about the market's long-run functional form. It substantially reduces dynamic bias at the cost of higher asymptotic variance at distant horizons. For causal-identification exercises centered on isolating the direct effect of a friction, this bias reduction makes Local Projections an appropriate methodological framework Jordà (2023).

4.2 Heterogeneity Control and Systemic Shocks

For the estimates to admit a strictly causal interpretation, the model must satisfy restrictive assumptions regarding the behavior of the error term $\epsilon_{i,t+h}$. Under classical Ordinary Least Squares (OLS), errors are assumed to be spherical: they have constant variance (homoskedasticity) and are strictly independent of one another. Formally, this implies that the covariance between the error for product i and product j at time t is zero:

$$\text{Cov}(\epsilon_{i,t}, \epsilon_{j,t}) = 0 \quad \text{para todo } i \neq j \tag{11}$$

In high-frequency e-commerce micro-panels, however, this assumption of spatial or cross-sectional independence is systematically violated.

4.2.1 The Problem of Cross-Sectional Dependence (CSD) and Heterogeneity

On a centralized *marketplace*, sellers do not operate in isolation. When a systemic macroeconomic event occurs —for example, a nationwide trucker strike or a sharp increase in fuel prices— thousands of merchants experience the disturbance simultaneously.

This interconnectedness generates what panel econometrics calls *Cross-Sectional Dependence* (CSD). The literature models CSD by allowing idiosyncratic errors to contain an unobserved multifactor structure rather than being purely random. The error term can be decomposed as follows:

$$\epsilon_{i,t} = \lambda_i' f_t + v_{i,t} \tag{12}$$

Where:

- f_t is a vector of unobserved common factors. In this setting, it represents the magnitude of the national logistics *shock* at time t.
- λ_i denotes heterogeneous factor loadings, that is, the product-specific sensitivity of item i to the *shock*. A heavy piece of furniture will have a much larger λ_i loading during a transportation strike than a lightweight technology product.
- $v_{i,t}$ is the true idiosyncratic error, or white noise, which is independent across sellers.

If the researcher ignores $\lambda_i' f_t$ and estimates the model using conventional OLS, a serious statistical problem arises. Because all sellers are exposed to the same factor f_t, their errors move in the same direction. The model then incorrectly treats thousands of observations as independent confirmations of a decline in demand when, in fact, they reflect a single systemic *shock* affecting thousands of products simultaneously. The result is a severe downward bias in the estimated residual variance, artificially inflated t-statistics, and an increased likelihood of *false discoveries.*

The standard method for addressing this form of spatial dependence and adjusting standard errors is the Driscoll-Kraay (1998) variance-covariance estimator. Constructing this matrix, however, requires evaluating the covariance between every pair of individuals in each period, implying algorithmic complexity of $\mathcal{O}(T \cdot N^2)$. In very large micro-panels, this calculation exhausts available memory and prevents numerical convergence, making estimation infeasible.

4.2.2 The Structural Solution: Two-Way Fixed Effects (TWFE)

Given this computational constraint, analytical econometrics offers a more efficient alternative. If correcting the error *ex-post* in the variance matrix is infeasible, the dependence can instead be removed *ex-ante* by modifying the model specification to include *Two-Way Fixed Effects* (TWFE).

The Local Projections model includes two sets of specific intercepts:

- **Individual Fixed Effect (α_i):** This term captures and absorbs time-invariant unobserved heterogeneity. SKU characteristics such as package volume or brand reputation can make a product's baseline sales permanently higher than those of its competitors. Including α_i effectively subtracts each individual's time mean, so identification comes only from deviations relative to that unit's own historical norm (*within variation*).
- **Time Fixed Effect (γ_t):** This component is central to neutralizing Cross-Sectional Dependence (CSD). Pesaran (2021) develops the argument that time fixed effects operate as an orthogonal projection that deterministically filters out systemic dependence.

To see how the time fixed effect neutralizes this dependence, rewrite the empirical equation by decomposing the factor loading λ_i. Let each individual's sensitivity equal the group-average sensitivity ($\bar{\lambda}$) plus an idiosyncratic deviation (η_i), so that $\lambda_i = \bar{\lambda} + \eta_i$. The contaminated error term can then be written as:

$$\lambda_i' f_t + v_{i,t} = (\bar{\lambda}' f_t) + (\eta_i' f_t) + v_{i,t} \tag{13}$$

This is where the physical-property clustering described in Section 3 (light goods versus heavy goods) acquires its econometric role. By grouping products with similar volumetric structures and logistics-pricing schedules, the design empirically compresses the dispersion in sensitivities toward zero ($\eta_i \approx 0$). It is therefore reasonable to treat products within the same physical cluster as having nearly identical exposure to the systemic event ($\lambda_i \approx \bar{\lambda}$).

Under this within-cluster homogeneity condition, saturating the model with the intercept γ_t removes the unobserved *shock* vector f_t from the residual variation because the fixed effect algebraically absorbs the corresponding cross-sectional component:

$$\gamma_t = \bar{\lambda}' f_t \tag{14}$$

The residual entering the estimation is then reduced to white noise: $\epsilon^*_{i,t} = v_{i,t}$. Algebraically, the time fixed effect absorbs the platform-wide average daily macroeconomic volatility. Once this strong Cross-Sectional Dependence (CSD) is purged by design, the remaining cross-unit spatial correlation falls from strong to weak and becomes asymptotically negligible for convergence of the coefficients of interest.

4.2.3 Serial Correlation and the Newey-West (HAC) Correction

Controlling for cross-sectional dependence with Two-Way Fixed Effects cleans the design matrix in the cross-sectional dimension. A final statistical issue nevertheless remains in the time dimension.

By construction, Local Projections mechanically induce serial correlation in the residuals. Because the model estimates the response at a future horizon $t+h$ using a *shock* observed at time t, the error term accumulates all new unobserved innovations affecting the product in the intervening periods $(t+1, t+2, \ldots, t+h)$.

Formally, the projected residual follows an overlapping *Moving Average* structure of order $h-1$, denoted MA$(h-1)$:

$$\epsilon_{i,t+h} = v_{i,t+h} + \psi_1 v_{i,t+h-1} + \cdots + \psi_{h-1} v_{i,t+1} \tag{15}$$

This structure implies that errors are not independent over time; analytically, $\text{Cov}(\epsilon_{i,t+h}, \epsilon_{i,t+h-1}) \neq 0$. Thus, if the model makes a forecast error on day 2, the error on day 3 will be correlated with it both empirically and mathematically.

To correct this temporal dependence and restore valid confidence intervals without changing the point estimates, the model uses one-dimensional heteroskedasticity-and-autocorrelation-consistent (HAC) standard errors based on the classical Newey-West Bartlett kernel.

The combination of TWFE and Newey-West is also supported by recent work. Abadie et al. (2023) show that imposing complex adjustments can produce excessively inflated and inefficient covariance matrices when the empirical design has already neutralized the primary source of aggregation. Because the systemic cross-sectional component is effectively isolated by the two-way structure, Newey-West standard errors provide the asymptotically appropriate correction, restoring valid inference without creating prohibitive computational burdens.

4.3 Treatment of Empirical Zeros and Scale Dependence

Once Cross-Sectional Dependence has been neutralized, the econometric design must address a structural feature of the dependent variable itself. High-frequency microdata contain a large mass of exact zeros. It is common for a particular product to record zero sales on a given day, whether because of weak demand or logistics frictions.

Classical elasticity estimation relies on logarithmic transformations. The structure of these data, however, prevents use of the standard log transformation because of its fundamental mathematical limit:

$$\lim_{y \to 0} \ln(y) = -\infty \tag{16}$$

Dropping zero-valued observations is not a methodologically acceptable solution. Excluding zero-sales days would introduce selection bias and destroy the panel structure.

4.3.1 The Conventional Solution: The Inverse Hyperbolic Sine (IHS)

To retain empirical zeros while preserving an elasticity-like interpretation, applied econometrics has often used the Inverse Hyperbolic Sine (IHS) transformation. The function is defined over the full real line and provides a continuous approximation to the natural logarithm:

$$\text{IHS}(y) = \ln\left(y + \sqrt{y^2 + 1}\right) \tag{17}$$

The most useful asymptotic property of this transformation is $\text{IHS}(0) = 0$, which preserves zero-sales days. For positive values ($y \geq 2$), the IHS converges to a shifted logarithmic form: $\text{IHS}(y) \approx \ln(2) + \ln(y)$.

Because of this convergence, the traditional literature often treated coefficients from log-linear and IHS-linear specifications as interchangeable standard elasticities. To see why that interpretation is not mathematically exact, consider the first derivative of the transformation with respect to the original variable y:

$$\frac{\partial \text{IHS}(y)}{\partial y} = \frac{1}{y + \sqrt{y^2 + 1}} \cdot \left(1 + \frac{1}{2\sqrt{y^2 + 1}} \cdot 2y\right) \tag{18}$$

Algebraically simplifying the terms inside parentheses reduces the derivative to:

$$\frac{\partial \text{IHS}(y)}{\partial y} = \frac{1}{\sqrt{y^2 + 1}} \tag{19}$$

For a linear model of the form $\text{IHS}(y) = \beta x$, applying the chain rule to recover the marginal effect of x on the original level variable y yields:

$$\frac{\partial y}{\partial x} = \beta \cdot \left(\frac{\partial \text{IHS}(y)}{\partial y}\right)^{-1} = \beta \cdot \sqrt{y^2 + 1} \tag{20}$$

As shown by Bellemare and Wichman (2020), this exact derivation demonstrates that marginal effects are not constant over the support of demand. The nonlinear adjustment term ($\sqrt{y^2 + 1}$) implies that the absolute magnitude of the response to a logistics *shock* is necessarily conditioned by the product's pre-shock sales level (y), in contrast to the constant elasticity implied by the conventional logarithmic model.

4.3.2 The Extensive-Margin Problem (Chen & Roth, 2024)

Beyond the nonlinear derivative correction, **?** document important structural limitations in estimating the Average Treatment Effect (ATE) when the outcome distribution contains a large mass of zeros, a problem formalized as *scale dependence*.

Their analysis shows that estimating a causal effect on an IHS-transformed variable generates results that depend on the units in which the variable is measured. Understanding the mechanism requires separating sellers' responses to a logistics *shock* into two margins:

- **Intensive Margin:** The change in sales volume conditional on the product remaining commercially active ($y > 0$).
- **Extensive Margin:** The marginal probability that the product records at least one transaction rather than no demand at all ($y = 0$).

If the unit of measurement is rescaled by multiplying the original variable by a factor $a > 0$ (for example, converting revenue from Brazilian reais to cents), the IHS transformation shifts the intensive-margin scale additively for sufficiently large positive values:

$$\text{IHS}(ay) \approx \ln(2ay) = \ln(2y) + \ln(a) \approx \text{IHS}(y) + \ln(a) \tag{21}$$

As the algebra shows, rescaling the intensive margin shifts all positive observations by the constant $\ln(a)$. At the extensive margin, however, when the logistics friction generates an empirical zero-sales day, the scalar collapses mechanically:

$$\text{IHS}(a \cdot 0) = \text{IHS}(0) = 0 \tag{22}$$

This mathematical asymmetry is the source of the structural bias. Rescaling shifts all positive sales observations by $\ln(a)$ while leaving zero-sales days rigidly anchored at the origin. The relative distance between a high-demand day and a zero-transaction day therefore changes mechanically. As a result, the magnitude of the estimated coefficient β_h depends on the scale of the panel, compromising the direct interpretation of the impulse response as a pure elasticity.

4.3.3 Academic Transparency and Computational Constraints

The technical solution to scale dependence is to abandon log-like transformations and use structural count-data models such as the Poisson Pseudo-Maximum Likelihood (PPML) estimator. PPML is mathematically invariant to rescaling and accommodates zeros directly.

In the present framework, however, implementation is computationally infeasible. PPML is subject to the *incidental parameters problem*[4] in settings that require the absorption of multiple spatial and time fixed effects (TWFE). Unlike linear models —where a product fixed effect can be removed analytically by subtracting its time average— nonlinear PPML must iteratively estimate an individual parameter for each of the network's thousands of SKUs. In high-frequency panels, this computational burden not only exhausts the matrix memory needed to estimate Local Projections horizon by horizon, but also allows estimation error in thousands of intercepts to contaminate and asymptotically bias the main coefficient.

For academic transparency, the analysis explicitly acknowledges the formal limitation identified by **?**. The use of the IHS transformation is driven by the need for a linear computational framework. To mitigate scale dependence, the monetary unit (Brazilian real, BRL) is held fixed across all estimates.

The demand-response coefficients (β_h) therefore should not be interpreted as pure structural semi-elasticities. Instead, they are composite log-like indices that combine absolute changes in *units sold* (the intensive margin) with changes in the marginal probability of conversion (the extensive margin). Empirically, a negative coefficient captures two features of the *marketplace* simultaneously: a reduction in purchase size (e.g., an SKU falling from 10 units per day to 5) and a transition to no transaction at all (e.g., an SKU falling from 1 unit to zero sales on that day because logistics costs increased).

This interpretive qualification preserves the causal direction, dynamic path, and relative comparability of effects across clusters while avoiding misleading structural-elasticity

[4] Count-data models are designed specifically for discrete, nonnegative dependent variables (such as the integer number of items sold in a day). They nevertheless face this classical econometric problem: in nonlinear maximum-likelihood models, the fixed effects cannot be separated algebraically from the rest of the equation (there is no exact *Within* transformation analogous to OLS). Requiring the algorithm to estimate one parameter for every product exhausts the sample's degrees of freedom and makes the estimator inconsistent unless the time dimension (T) tends to infinity.

interpretations.

4.4 Robustness to Unit Roots and Extreme Persistence

The final structural challenge concerns the stochastic properties of the time series. Classical econometric theory requires model variables to be stationary, or integrated of order zero, $I(0)$, for standard statistical inference to be valid.

Nominal prices in digital-market microstructure, however, follow a different pattern. Independent sellers adjust prices infrequently because of strict platform algorithms. This nominal rigidity produces stepwise series: prices remain unchanged for weeks, shift discretely, and then become flat again.

Statistically, this inertial behavior produces highly persistent series that often contain a Unit Root or are integrated of order one, $I(1)$. In the presence of $I(1)$ variables, the conventional approach repeatedly applies first differences ($\Delta P_t = P_t - P_{t-1}$) to force the series to become stationary before estimation.

4.4.1 The Risk of Destructive Overdifferencing

Although first differencing is the classical standard response, applying it to this micropanel would create a serious methodological limitation. First, differencing removes the structural information contained in the absolute levels of the variables.

From a behavioral-economics perspective, consumers do not respond only to the daily price change (ΔP_t); they respond to the absolute *Landed Cost*. Freight aversion arises because consumers compare the absolute freight charge with the product's base price. Differencing prices would therefore remove this fundamental cognitive information —the ratio that anchors the consumer's decision.

Second, from a purely statistical perspective, differencing very high-frequency data tends to induce overdifferencing and mechanically inflate variance. Suppose observed daily prices contain purely random measurement white noise ε_t with constant variance σ^2. Taking the first difference ($\Delta\varepsilon_t = \varepsilon_t - \varepsilon_{t-1}$) doubles the variance of the resulting error:

$$\mathrm{Var}(\varepsilon_t - \varepsilon_{t-1}) = \mathrm{Var}(\varepsilon_t) + \mathrm{Var}(\varepsilon_{t-1}) = 2\sigma^2 \tag{23}$$

Accordingly, the strict differencing required by conventional econometrics would double the statistical noise and obscure the underlying economic signal of the logistics effect.

4.4.2 The Analytical Solution: *Lag-Augmented Local Projections*

To reconcile valid inference with preservation of variables in levels, the analysis relies on the results developed by Montiel Olea and Plagborg-Møller (2021).

The authors show analytically that Local Projection inference can robustly accommodate extreme persistence, making overdifferencing unnecessary. Their central contribution is the formalization of *Lag-augmented Local Projections*.

The econometric mechanism relies on the Frisch-Waugh-Lovell (FWL) theorem. Suppose the endogenous variable (for example, price y_t) follows a highly persistent autoregressive Data-Generating Process (DGP):

$$y_t = \rho y_{t-1} + u_t \tag{24}$$

Where $\rho \approx 1$ or $\rho = 1$ (indicating a Unit Root) and u_t is an exogenous innovation, that is, the purely random *shock* or "surprise" realized on that day.

If the researcher estimates a standard local projection without controls, the regression inherits the full memory and nonstandard statistical properties of a *random walk*. Montiel Olea and Plagborg-Møller (2021) show, however, that proactively including lags of the dependent and independent variables in the specification (represented by the sums $\sum_{j=1}^{p} \delta_{j,h} y_{i,t-j} + \sum_{j=1}^{p} \theta_{j,h} C_{i,t-j}$) causes the Ordinary Least Squares (OLS) algorithm to purge this persistence automatically.

The intuition does not require complex matrix algebra. Including yesterday's price (y_{t-1}) as a control effectively instructs the model to "estimate today's logistics effect while holding constant yesterday's price inertia."

The algorithm then projects the regressor onto its own historical lags and removes the predictable inertial component. The model no longer estimates the elasticity using a series contaminated by infinite memory, but rather a residualized, or purged, regressor ($\tilde{y}_t$), which is effectively the difference between what occurred today and what was known yesterday:

$$\tilde{y}_t \approx y_t - \rho y_{t-1} = u_t \tag{25}$$

This result is central: the residualized variable ($\tilde{y}_t$) is precisely the stochastic innovation (u_t). Because u_t captures only that day's temporal "surprise" and carries no inherited history, it is by definition a stationary $I(0)$ process.

Consequently, once lags are included as controls, the Local Projection estimator β_h converges asymptotically to a standard Normal distribution. This ensures valid test statistics and confidence bands for the Impulse Response Functions even though the data have not been manually differenced.

On this mathematical basis, the econometric model is estimated directly in nominal levels: log price $\ln(P_t)$ and IHS-transformed demand $IHS(Q_t)$.

Rather than differencing these series, the algorithm parametrically includes a lag vector of order $p = 2$ for both logistics shocks and historical endogenous variables. The use of two temporal lags (48 hours) is intended to absorb not only the immediate inertia from the preceding day but also short operational cycles typical of e-commerce (such as weekend warehouse-processing delays). In this way, Lag-augmentation isolates the genuinely exogenous logistics shock, provides the dynamic estimates with correct asymptotic coverage, and protects the results against conventional stationarity concerns.

5 DIAGNOSTICS AND STATISTICAL PROPERTIES

The theoretical foundations developed in the methodological section impose a set of stochastic assumptions on the data. In applied econometrics, however, model validity cannot rest on theory alone; these assumptions must be evaluated using empirical diagnostics on the observed data matrix.

This section diagnoses the statistical properties of the micro-panel along three dimensions that are central to the consistency of the Local Projections model: cross-unit dependence (the Pesaran test), exogeneity of unobserved effects (the Hausman test), and persistence or stationarity of the time series (the Fisher-Maddala-Wu Unit Root test). These diagnostics assess the parametric choices adopted above and reveal structural features of price-setting dynamics on digital platforms.

5.1 Cross-Sectional Dependence Test (Pesaran CD)

As argued in the model specification, the logistics infrastructure suggests that sellers are exposed to contemporaneous systemic shocks. Changes in the national logistics operator's tariffs or supply-chain inefficiencies do not affect products in isolation; they propagate across the network.

To test whether this macroeconomic disturbance induces spurious correlation across product-level residuals —thereby violating the classical cross-sectional independence condition $Cov(\epsilon_{i,t}, \epsilon_{j,t}) = 0$— the analysis implements the Cross-Sectional Dependence (CD) test developed by Pesaran (2021).

This test is designed specifically for panels with a large cross-sectional dimension (N). Unlike the classical Breusch-Pagan Lagrange Multiplier test, which becomes distorted and inconsistent when N is large, Pesaran's statistic directly evaluates the scaled sum of pairwise sample residual correlations ($\hat{\rho}_{ij}$) obtained after Fixed Effects estimation. Formally, the test statistic is:

$$CD = \sqrt{\frac{2T}{N(N-1)}} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \hat{\rho}_{ij}$$

Under this formulation, the leading scaling factor normalizes the sum of cross-correlations so that, under the null hypothesis, the statistic converges asymptotically to a standard Normal distribution ($N(0,1)$)[5]. The hypotheses are:

- **Null Hypothesis (H_0):** The errors are cross-sectionally independent (no CSD; $\hat{\rho}_{ij} = 0$).
- **Alternative Hypothesis (H_1):** The errors exhibit contemporaneous Cross-Sectional Dependence.

Empirical Results and Methodological Implications:

The Pesaran test was run separately for the two outcome equations and for the two physical partitions of the dataset (Light Cluster and Heavy Cluster), yielding the following z-statistics:

- **Cluster 1 (Light Goods):**
  - Price Model: $z = 27.022$, p-value $< 2.2 \times 10^{-16}$.
  - Demand Model: $z = 31.530$, p-value $< 2.2 \times 10^{-16}$.
- **Cluster 2 (Heavy and Bulky Goods):**
  - Price Model: $z = 16.031$, p-value $< 2.2 \times 10^{-16}$.
  - Demand Model: $z = 25.071$, p-value $< 2.2 \times 10^{-16}$.

Because the p-values converge effectively to zero across all catalog segments, the null of independence is decisively rejected. The empirical results confirm strong Cross-Sectional Dependence (CSD) in the panel regardless of the physical characteristics of the good.

This result is economically and statistically important because it shows that estimating a

[5] Intuitively, the test takes the model's prediction "error" (the residual) and pairs products two at a time (e.g., product A with product B) to assess whether their errors rise or fall together. Summing correlations across all possible pairs in the network produces a very large and difficult-to-interpret number. The square-root term acts as a "mathematical normalizer": it rescales this sum using the number of products (N) and days (T) so that it maps onto a standard scale (a Gaussian distribution). A value of 0 therefore corresponds to independence, whereas an extreme value signals that the market moves jointly.

pooled Ordinary Least Squares (OLS) model would produce invalid inference by understating the true standard error. It also supports the parametric choice developed in the methodology section: CSD is sufficiently strong that time fixed effects within the TWFE specification are structurally necessary to remove the unobserved common factor before computing the Impulse Response Functions with Newey-West HAC errors.

5.2 Hausman Test

One of the central decisions in panel-data econometrics concerns the treatment of individual unobserved heterogeneity, represented parametrically by α_i. In this application, α_i captures product characteristics that are constant over time but unobserved in the data, such as intrinsic brand reputation, the seller's operational efficiency, or the geographic location of its primary warehouse.

The literature offers two competing parametric approaches: the *Random Effects* (RE) model and the *Within*, or *Fixed Effects* (FE), estimator. The choice hinges on strict exogeneity. RE assumes that individual unobserved heterogeneity is uncorrelated with the explanatory variables: $\text{Cov}(\alpha_i, C_{i,t}) = 0$. If this condition holds, RE is the more efficient estimator.

If, by contrast, α_i is correlated with logistics freight ($\text{Cov}(\alpha_i, C_{i,t}) \neq 0$), the RE estimator suffers from severe omitted-variable bias and becomes inconsistent. In that case, FE is the only asymptotically valid estimator because the Within transformation removes α_i from the estimating equation by subtracting the unit-specific time mean.

To adjudicate between these alternatives formally, the analysis implements the specification test of Hausman (1978). The statistic measures the "distance" between the coefficients estimated under the two models, weighted by the difference between their variance matrices. Formally:

$$H = (\hat{\beta}_{FE} - \hat{\beta}_{RE})'[\hat{V}_{FE} - \hat{V}_{RE}]^{-1}(\hat{\beta}_{FE} - \hat{\beta}_{RE}) \tag{26}$$

Under the null, this statistic is asymptotically distributed as Chi-square (χ^2_k) with k degrees of freedom (equal to the number of time-varying regressors). The hypotheses are:

- **Null Hypothesis (H_0):** The unobserved effects are uncorrelated with the regressors ($\text{Cov}(\alpha_i, X_{i,t}) = 0$). Both estimators are consistent, but RE is asymptotically more efficient.
- **Alternative Hypothesis (H_1):** The unobserved effects are correlated with the regressors. The Random Effects model is structurally inconsistent, making the Fixed Effects estimator necessary.

Empirical Results and Causal Justification:

The model was tested by estimating both specifications separately for each cluster. The Hausman test produced the following results:

- **Cluster 1 (Light Goods):** $\chi^2 = 329.52$, $df = 2$, p-value $< 2.2 \times 10^{-16}$.
- **Cluster 2 (Heavy Goods):** $\chi^2 = 45.08$, $df = 2$, p-value $= 1.62 \times 10^{-10}$.

Because the p-value strongly rejects the null in both cohorts, the empirical evidence indicates that the Random Effects estimator is structurally inconsistent[6].

From an economic and organizational perspective, this result is consistent with the institutional setting. Unobserved seller characteristics (such as the geographic location of a logistics center or warehousing agreements) determine distance to the average consumer and therefore the seller's baseline freight-cost structure. Assuming that unobserved product identity is unrelated to the cost structure would generate substantial endogeneity. The Hausman test therefore provides both statistical and causal support for using TWFE estimators in the Local Projections specification.

5.3 Stationarity Test (Fisher-Maddala-Wu)

The final diagnostic concerns the order of integration of the time series. It is important to determine whether the variables are stationary —integrated of order zero, $I(0)$— or instead contain a Unit Root that makes them highly persistent and nonstationary — integrated of order one, $I(1)$. An $I(0)$ series tends to revert toward its historical mean following a shock. An $I(1)$ series, by contrast, has infinite memory: shocks have permanent effects over time, as in a random walk.

To evaluate stationarity, the data design uses dynamic Lifecycle Bounding, producing a markedly unbalanced panel in which each SKU has a different time dimension T_i. In this setting, the Fisher-Maddala-Wu (1999) test is particularly suitable because it relaxes the requirement that the panel be balanced.

This nonparametric procedure estimates an Augmented Dickey-Fuller (ADF) test separately for each product i and extracts its p-value (p_i). The appeal of Fisher's test follows from the uniform distribution: under the null, continuous p-values are uniformly distributed between 0 and 1. Hence, the transformation $-2\ln(p_i)$ follows a Chi-square distribution with 2 degrees of freedom. Summing across N independent variables produces

[6] The reported statistics correspond to the classical Hausman test, which assumes homoskedasticity under the null. Transaction-level micro-panels commonly exhibit structural heteroskedasticity, which would in principle motivate a variance-covariance-estimator (VCE) robust variant. The empirical rejection is nevertheless so strong (with asymptotically negligible p-values) that use of the standard test does not alter the inferential conclusion, making the choice of the Fixed Effects estimator effectively unavoidable.

the global statistic (MW), which is asymptotically distributed as χ^2 with $2N$ degrees of freedom[7]:

$$MW = -2\sum_{i=1}^{N} \ln(p_i) \sim \chi^2_{2N}$$

The hypotheses are defined as follows:

- **Null Hypothesis (H_0):** All time series in the panel contain a Unit Root (they are $I(1)$).
- **Alternative Hypothesis (H_1):** At least a nontrivial fraction of the panel series is stationary ($I(0)$).

Empirical Diagnosis and the Zero-Variance Problem:

When the algorithm was first applied to nominal log prices ($\ln P_t$), the test returned missing values (NaN). Rather than indicating a programming failure, this result revealed a structural feature of the ecosystem: extreme algorithmic rigidity. A large share of SKUs maintained an exactly constant price throughout their observed lifetime, implying zero variance ($\sigma^2 = 0$).

Formally, the ADF statistic divides the estimated coefficient by its standard error. If the variable is an immutable constant, the variance-covariance matrix becomes singular and the standard error collapses to zero, making the ratio undefined. To resolve this indeterminacy and obtain valid inference, the implementation applies a screening rule and computes the test only for SKUs with more than 10 days of observed life and strictly positive variance ($var > 0$).

Under this restriction, the final results for the three core variables are:

- **Demand ($IHS(Q_t)$):** p-values of 7.80×10^{-78} (Light) and 9.22×10^{-45} (Heavy). H_0 is rejected; the variable is stationary, $I(0)$.
- **Logistics Cost ($IHS(C_t)$):** p-values of 2.47×10^{-76} (Light) and 1.60×10^{-38} (Heavy). H_0 is rejected; the variable is stationary, $I(0)$.
- **Selling Price ($\ln P_t$):**
  - Light Cluster ($N = 90$ active): $MW = 180.44$, p-value $= 0.4766$.

[7] Intuitively, Fisher's test acts as an "aggregator of evidence." If a single product yields a p-value of 0.15, there is insufficient evidence to classify it as stationary. If, however, 50 products all yield the same value of 0.15, the probability of observing that pattern by chance becomes extremely small. The natural logarithm (ln) maps small p-values into large negative numbers. Multiplying by -2 and summing causes the global statistic to increase sharply, allowing the null to be rejected jointly for the platform even when individual series are too short or noisy to be conclusive on their own.

– Heavy Cluster ($N = 50$ active): $MW = 104.41$, p-value $= 0.3616$.

Economic and Methodological Implications:

The results are stark. For demand and freight cost, the extremely small p-values reject the null and confirm that both series are $I(0)$. In other words, logistics fluctuations are transitory.

For the base-price variable ($\ln P_t$), however, the test clearly fails to reject the null. Empirically, this indicates that prices contain a Unit Root ($I(1)$). From an Industrial Organization perspective, the result is consistent with the extreme nominal rigidity governing the platform: prices follow an inertial stepwise process under algorithmic penalties and do not dynamically revert to their pre-shock mean.

Methodologically, this finding supports the architecture developed in Section 4. Given the statistical evidence that price is $I(1)$, a conventional VAR would require differencing, thereby removing the absolute price-level information that anchors consumer decisions. The model therefore implements *Lag-augmented Local Projections* (Montiel Olea and Plagborg-Møller (2021)) in levels, providing an inferential framework designed to remain valid under extreme persistence.

6 RESULTS: STRUCTURAL DYNAMICS AND ASYMMETRIC IMPACTS

The empirical validation of the econometric assumptions discussed in Section 5 —in particular, the appropriateness of level estimators in the presence of variables integrated of order one, $I(1)$— provides the statistical basis for extracting and interpreting the causal estimates. The analysis now moves from statistical validation of the sample to applied economic inference.

The central objective of this chapter is to quantify the *cost pass-through* mechanism within the Olist *marketplace*. Specifically, the analysis evaluates how sellers adjust their pricing strategies and how aggregate demand responds when the platform experiences an exogenous *shock* to logistics charges. To capture the ecosystem's intrinsic heterogeneity and avoid aggregation bias, the empirical results separate the market into the two physical environments defined in the data strategy: the Light Goods cluster (technology, health, beauty) and the Heavy and Bulky Goods cluster (furniture, tools, appliances).

6.1 General Empirical Framework: Anatomy of a Logistics Shock

The econometric device used to visualize and quantify these dynamic effects is the Impulse Response Function (IRF). The IRFs are constructed directly from the sequence of β_h coefficients estimated by the system of Local Projections (LP) specified in Equation (10).

Before turning to the behavior of each submarket, it is useful to establish a precise framework for interpreting these graphical projections in a high-frequency micro-panel. The IRFs are read according to the following statistical principles:

- **Time Horizon (Horizontal Axis):** The horizontal axis reports the projection horizon, h, measured in calendar days. The model traces the dynamic effect from the moment the logistics cost increase occurs ($t = 0$) through a forward horizon of fifteen days ($t = 15$). This horizon is long enough to capture both the initial response and the asymptotic stabilization of pricing algorithms.
- **Magnitude of the Effect (Vertical Axis):** The vertical axis measures the response of the endogenous variables (nominal price $\ln P_t$ or demand $\text{IHS}(Q_t)$) to a positive, orthogonalized exogenous impulse equal to one standard deviation in freight cost (C_t).
- **No Causal Effect (Zero Line):** The horizontal line at $y = 0$ represents the no-effect benchmark. For an effect to be statistically significant, the point estimate (the solid line in the figure) must depart from this boundary.
- **Uncertainty and Confidence Bands:** In observational econometrics, each point estimate must be accompanied by a measure of uncertainty. In the projections, the shaded regions delimit the confidence intervals for β_h. Both 90% and 95% confidence bands are reported. Their width is calibrated using Newey-West variance estimators combined with Two-Way Fixed Effects (TWFE), which purge contemporaneous Cross-Sectional Dependence from the panel. The inferential rule is straightforward: if, at horizon h, a confidence band intersects $y = 0$, the causal effect of the *shock* is statistically indistinguishable from zero.

Under this interpretation, the estimates show that e-commerce does not respond mechanically to external inflationary pressures. Subjecting the full panel to an identical cost *shock* reveals that nominal rigidity in the base price is systemic, while the demand consequences differ sharply with the financial viability and physical characteristics of the good.

6.2 Cluster 1 (Light Goods): Nominal Rigidity and Algorithmic Coercion

The first cluster, comprising light parcel goods and impulse-purchase items (technology, beauty, watches), yields evidence that challenges standard predictions of digital microeconomics. The Impulse Response Function for the base price ($\ln P_t$) indicates extreme nominal rigidity. At the exact moment the shipping-cost *shock* occurs ($t = 0$), the estimated price response is essentially zero: $\beta_0 = -0.0006$. More importantly, the confidence interval comfortably includes zero ($\text{IC} : [-0.0022, 0.0010]$), providing no statistical evidence of a causal effect.

This parametric immobility is not merely a short-run delay in algorithmic adjustment. Across all fifteen horizons, the estimated path remains statistically indistinguishable from zero. At the end of the projection window ($t = 14$), the estimate is $\beta_{14} = -0.0033$, with a positive upper confidence bound (0.0003), again indicating that the seller does not pass the freight increase through to the final consumer.

Standard economic theory would predict that digitalization, by reducing *menu costs*, should generate a highly flexible environment in which prices adjust continuously. The data show the opposite pattern, consistent with a recent strand of Industrial Organization research that emphasizes the restrictive role of platform algorithms.

6.2.1 Centripetal Technology and *Buy Box* Governance

The lack of price adjustment is consistent with the consolidation of *Algorithmic Pricing*. Brown and MacKay (2025) characterize artificial-intelligence-based *pricing* in *marketplaces* as a "centripetal technology." Unlike centrifugal forces that promote dispersion, algorithms pull prices toward a common template and penalize unilateral upward deviations.

The instrument enforcing this market discipline is the *Buy Box*, an algorithmic position through which the overwhelming majority of conversions are routed. Gleason et al. (2026) show empirically that when an algorithm detects a price increase above competitors' average —even when the increase reflects an exogenous cost shock— the platform may suppress the seller's purchase button through *Buy Box suppression*.

For a light-goods seller, the marginal consequence of raising price is therefore the loss of algorithmic visibility. Lee and Musolff (2025) quantify this penalty as equivalent to a 16% price surcharge relative to internal platform competitors. The rational seller consequently keeps the base price anchored (captured by $\beta \approx 0$) to avoid this severe visibility penalty.

6.2.2 *Markup Compression*: The Seller as Shock Absorber

If the consumer price remains unchanged following a positive freight *shock*, accounting identities require another margin to absorb the cost. The model implies that this margin is the seller's operating contribution margin, a mechanism of *Markup Absorption*. Alvarez-Blaser et al. (2025) document that during logistics-inflation episodes, retailers reallocate margins internally, compressing their *markup* to protect the final posted price.

This strategy achieves its commercial and algorithmic objective. The Impulse Response Function for Demand ($\text{IHS}(Q_t)$) shows a large initial decline from the logistics friction at $t = 0$ ($\beta_0 = -0.3605$). Because the base price remains unchanged, however, demand stabilizes relatively quickly and, by the end of the horizon ($t = 14$), the decline is no longer statistically significant ($\beta_{14} = -0.0431$, with a confidence interval crossing zero: $[-0.0883, 0.0021]$).

Cluster 1 Conclusion: In the light-goods segment, the seller absorbs the cost in order to preserve sales volume under algorithmic pressure. The digital ecosystem therefore acts as a deflationary shield for the consumer at the expense of seller profitability.

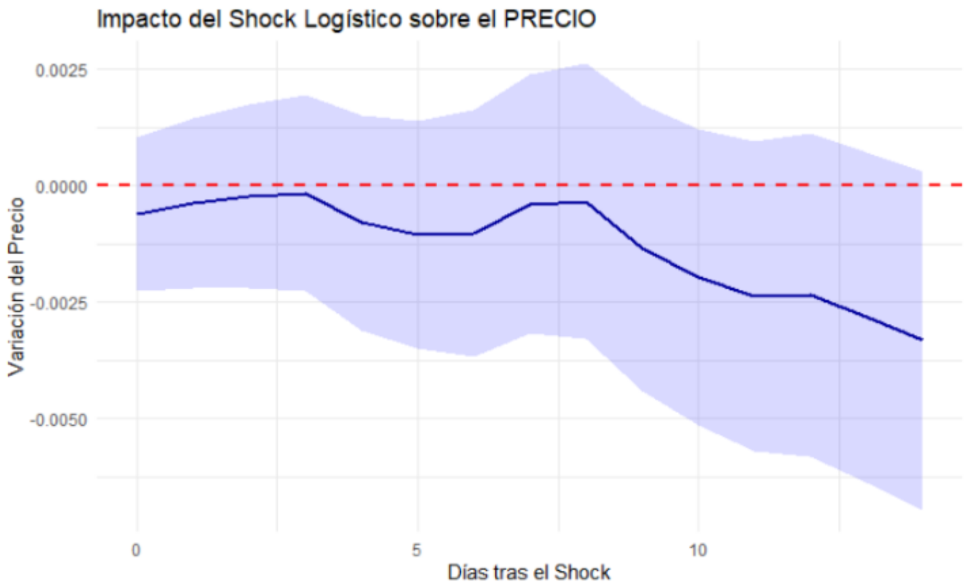

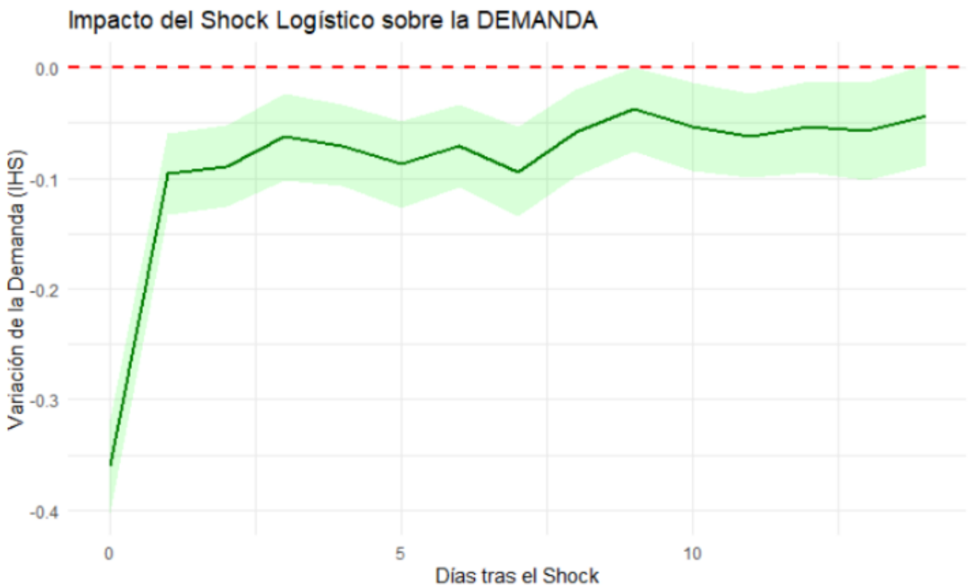

Figure 1: Impulse Response Functions (IRFs) for the Light Goods Cluster following a 1 S.D. *shock* to logistics cost. The base-price response remains essentially zero over the projection horizon, while the demand effect attenuates over time. Confidence bands are estimated at the 90% and 95% levels (TWFE + Newey-West).

6.3 Cluster 2 (Heavy Goods): Baseline Rigidity, Freight Salience, and Offline Flight

The analysis of Cluster 2 —heavy and bulky goods such as furniture and large appliances—reveals the same form of algorithmic rigidity across sellers but a much sharper deterioration in demand because of the logistics cost structure.

For the price-level IRF ($\ln P_t$), the results indicate base-price nominal rigidity similar to that observed for light goods. At the time of the logistics *shock* ($t = 0$), the estimated effect is effectively zero ($\beta_0 = -0.0004$), and the confidence interval contains zero (IC : $[-0.0019, 0.0012]$). Although the point estimate begins to trend slightly upward after the sixth day (reaching a maximum at $t = 11$ of $\beta_{11} = 0.0014$), the confidence interval contains zero at every one of the fifteen horizons (at $t = 11$, it ranges from -0.0009 to 0.0037). Statistically, the base-price restriction remains complete.

This result reinforces the algorithmic-coercion mechanism: price-parity policies do not vary with a product's physical volume or cost structure.

Unlike light goods, however, freight on heavy merchandise represents a critical share of Cost of Goods Sold (COGS). Sellers do not have sufficient operating margin to absorb the *shock* through *Markup Absorption*. Caught between the algorithmic restriction and the financial impossibility of absorbing the net loss, the *shock* is transmitted to the market through a larger *Partitioned Shipping Fee* charged explicitly to the consumer at *checkout*.

6.3.1 Demand Collapse: Aversion and Haptic Bias

The demand response to this increase in logistics cost is severe. The sales-volume IRF ($\text{IHS}(Q_t)$) shows an abrupt and statistically significant contraction, with an impact of $\beta_0 = -0.3303$ at $t = 0$ and a confidence interval entirely below zero (IC : $[-0.3793, -0.2812]$).

The key distinction relative to Cluster 1 is the complete absence of recovery over the observed horizon. At $t = 14$, the estimate remains firmly negative ($\beta_{14} = -0.0511$), and the upper confidence bound remains below zero (IC : $[-0.0937, -0.0085]$). Consumers do not merely postpone the purchase; for that transaction, they leave the digital ecosystem.

Recent research helps interpret this pattern through the interaction of behavioral economics and logistics inefficiencies:

- **Aversion to Partitioned Pricing:** Tran (2024) shows that buyers display strong aversion to partitioned shipping charges when these become salient and large at the end of the purchase funnel, sharply increasing cart abandonment.
- **Haptic Heaviness Bias:** Roy et al. (2024) document that the "haptic sensation of

heaviness" shapes *pricing* preferences. For bulky products, consumers show a strong preference for *combined pricing* and reject structures in which freight is charged separately.

6.3.2 Spatial Substitution: The Shift to Physical Retail (*Offline Flight*)

Because the need to purchase bulky goods is often relatively inelastic (for example, replacing a failed appliance), the demand that disappears from the parametric estimates is consistent with a direct shift toward traditional physical retail (*Offline Flight*).

? document that in *Online-Merge-Offline* (OMO) ecosystems, physical *retail* consolidates inventories through wholesale economies of scale and embeds logistics inefficiencies directly in shelf prices. By offering a consolidated price and removing the salience of last-mile freight, the physical channel presents the logistics surcharge in a less salient form, accommodates consumers' haptic bias, and recaptures market share from the digital channel.

Cluster 2 Conclusion: The econometric model indicates that the parametric rigidity imposed on the base price ($\beta \approx 0$) makes freight increases highly visible; interacting with consumer cognitive biases, this produces a persistent decline in digital demand and encourages substitution toward the physical channel.

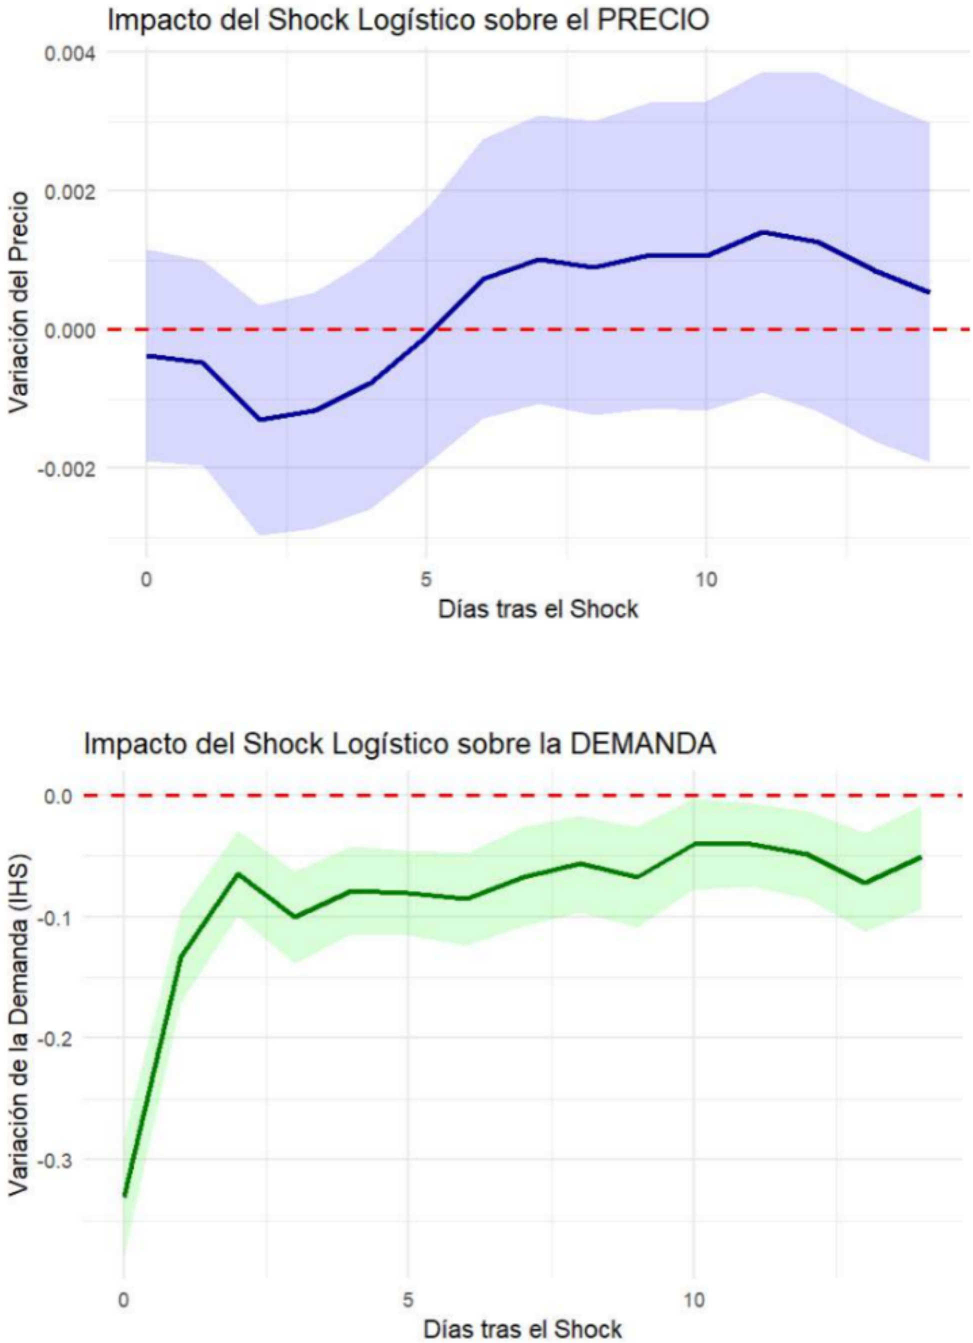

Figure 2: Impulse Response Functions (IRFs) for the Heavy and Bulky Goods Cluster. The base price remains persistently rigid (the confidence interval contains zero throughout the projection horizon), while demand contracts sharply and persistently (upper CI < 0 at $t = 14$). Confidence bands are estimated at the 90% and 95% levels (TWFE + Newey-West).

7 CONCLUSION AND CORPORATE IMPLICATIONS

7.1 Analytical Synthesis: Algorithmic Visibility vs. Operating Margin

This study set out to identify the causal transmission mechanisms of logistics *shocks* in high-frequency e-commerce environments and to assess whether sellers operate under the theoretical flexibility associated with digitalization or instead face systemic constraints. Using Local Projections (LP) estimated on an unbalanced longitudinal panel, the empirical findings challenge the classical logic of *cost-plus pricing* in platform markets.

The model's central finding is evidence of systemic algorithmic coercion. When the network is exposed to an exogenous increase in logistics freight, the nominal base price ($\ln P_t$) remains statistically unchanged ($\beta \approx 0$) across all estimated horizons, regardless of product volume or cost structure. This nominal rigidity indicates that in highly competitive ecosystems, short-run profit maximization is subordinated to algorithmic survival: sellers prioritize retaining *Buy Box* visibility over protecting gross operating margins.

The demand response, however, reveals a sharp divide in the viability of last-mile *e-commerce* that depends on the physical characteristics of the traded good:

- **Resilience through Compression (Light Cluster):** For low-dimensional-weight parcels, the seller absorbs the *shock* by reducing its own contribution margin through *Markup Absorption*. This mechanism acts as an implicit subsidy that shields the final consumer from the logistics friction, allowing the initial negative demand effect to attenuate and lose statistical significance at longer horizons.

- **Structural Contraction and *Offline Flight* (Heavy Cluster):** Heavy and bulky goods, by contrast, lack sufficient financial resilience. The inability to compress margins, combined with the algorithmic constraint on raising the base price, forces sellers to impose high partitioned freight charges at *checkout*. The model indicates that this visible surcharge interacts with consumer cognitive biases (aversion to logistics surcharges and haptic bias), producing an abrupt and persistent decline in digital demand. The contraction in sales volume is consistent with a direct reallocation of market share toward traditional physical retail (*Offline Flight*), where logistics costs are less salient because they are embedded in consolidated shelf prices.

In summary, this Master's thesis concludes that platform algorithms act as an effective deflationary shield for light parcel goods but become a mechanism of demand destruction for bulky merchandise when the market is exposed to exogenous logistics disturbances.

The following table provides a visual summary of the distinct causal paths and seller strategies observed in the two submarkets:

Table 1: Structural Summary of IRF Paths Following an Exogenous Logistics Shock

Endogenous Variable	Cluster 1: Light Goods	Cluster 2: Heavy Goods	Empirical Causal Interpretation
Base Price (P_t)	No change: the confidence interval contains $y = 0$.	No change: persistent rigidity; the interval contains $y = 0$.	Algorithmic Coercion: the platform penalizes changes in the base price.
Logistics Surcharge	Absorption of the *shock* through compression of the net margin.	Explicit pass-through to consumers through freight at *checkout*.	*Markup Compression* (Light) vs. Partitioned Pricing (Heavy).
Demand (Q_t)	Stability: attenuation of the effect over the medium run.	Sharp and persistent decline: no rebound over the horizon.	Volume retention (Light) vs. *Offline Flight* (Heavy).

7.2 Implications for Platforms: Towards a Resilient Market Design

The empirical findings indicate that the current pricing architecture of *marketplaces* is suboptimal under logistics stress. Platforms cannot assume that rigid parity algorithms maximize long-run ecosystem value. To protect Gross Merchandise Value (GMV) and limit the diversion of economic activity toward the physical channel, Operations and *Pricing Science* teams should consider three structural interventions:

1. **Context-Dependent Algorithmic Flexibility:** The *machine learning* systems governing the *Buy Box* should move from strict penalties toward conditional evaluation. If the platform detects a systemic exogenous *shock* in the logistics network, the algorithm should incorporate a "*pass-through* tolerance." This would allow sellers of bulky goods to incorporate the surcharge into the base price without suffering a visibility penalty, avoiding disproportionate freight charges at the end of the conversion funnel.

2. ***Cross-Subsidization* of Logistics:** Because heavy goods experience a sustained loss of demand under partitioned freight pricing, platforms should internalize this risk. The ecosystem could subsidize the last mile for bulky goods by financing the intervention through a marginal increase in the *take-rate* (transaction commission) on light goods. As the model indicates, seller-side markup absorption in the latter segment preserves sales volume without destroying final demand.

3. **Mandatory Consolidated-Pricing Architecture:** To reduce haptic bias and psychological aversion to partitioned surcharges in Cluster 2, the platform's *Market Design* should alter how prices are displayed. For heavy goods, the platform could contractually require or strongly incentivize a "Free Shipping" standard, forcing sellers to embed the true logistics cost into a unified posted price through *combined pricing*. This structure competes directly with physical *retail* shelf prices on comparable psychological terms.

7.3 Empirical Limitations and Future Lines of Research

Despite the robustness of the two-way panel estimators and the spatial variance corrections applied above, the study recognizes limitations arising from both the structure of the observational data and the mathematical properties of the dynamic estimators. These limitations open clear avenues for future research.

First, using the Inverse Hyperbolic Sine (IHS) transformation to accommodate empirical zeros in the Local Projections introduces the *scale dependence* documented by **?**. Although the comparative analysis identifies the direction of the effect, the estimated coefficient necessarily combines the intensive margin (the decline in the number of units sold) with the extensive margin (the probability that a seller temporarily records no activity). Future work with greater computational capacity should replicate the framework using Poisson Pseudo-Maximum Likelihood (PPML) estimators with iterative absorption of high-dimensional spatial fixed effects, thereby recovering a cleaner structural elasticity.

Second, a fundamental limitation of observational platform data is algorithmic opacity. The model infers the punitive effect of the *Buy Box* from inertial pricing behavior but does not observe the platform's internal auction algorithm. A causal validation of this mechanism would require a large-scale Randomized Controlled Trial (RCT or A/B Testing). A design in which the algorithm temporarily relaxes price-parity restrictions for a treatment group would allow the retention effect to be isolated relative to a control group operating under the standard rules.

Finally, a natural extension of this framework lies at the intersection of causal inference and *Causal Machine Learning*. One promising extension would replace the linear Local Projections regression with *Double Machine Learning* (DML). This would allow the analysis not only to control flexibly for highly nonlinear confounding but also to estimate Conditional Average Treatment Effects (CATE). From a business perspective, a DML model could reveal hidden heterogeneity across the network, identifying the geographic microregions or customer profiles for which sensitivity to partitioned freight generates the largest shift toward the physical channel and thereby improving the real-time allocation of logistics subsidies.

References

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2023). When Should You Adjust Standard Errors for Clustering? *The Quarterly Journal of Economics*, 138(1), 1-35.

Aguiar, V., & Kashaev, N. (2024). Random Utility Models with Unobserved Choice Sets. *Econometrica*, Forthcoming.

Alvarez-Blaser, [Inic.], Cavallo, A., MacKay, A., & Mengano, [Inic.] (2025). Markup Absorption and the Transmission of Logistics Shocks in E-commerce. *Working paper.*

Baymard Institute (2024). 50 Cart Abandonment Rate Statistics 2026 – Cart & Checkout. *Retrieved from baymard.com*

Bellemare, M. F., & Wichman, C. J. (2020). Elasticities and the Inverse Hyperbolic Sine Transformation. *Oxford Bulletin of Economics and Statistics*, 82(1), 50-61.

Brown, Z. Y., & MacKay, A. (2023). Competition in Pricing Algorithms. *American Economic Journal: Microeconomics*, 15(2), 109-56.

Brown, Z. Y., & MacKay, A. (2025). Algorithmic Pricing and Market Competition: The Centripetal Technology. *National Bureau of Economic Research (NBER) Working Paper.*

Cargas e Transportes (2022). Desperdício de combustível. *Retrieved from cargasetransportes.com.br*

Cavallo, A. (2018). Scraped Data and Sticky Prices. *Review of Economics and Statistics*, 100(1), 105–119.

Chen, J., & Roth, J. (2024). Logs with Zeros? Some Problems and Solutions. *The Quarterly Journal of Economics*, 139(2), 891-936.

Chen, Y., & Roth, J. (2024). *Log-like transformations and zeros in empirical research.* Journal of Econometrics.

Chudik, A., & Pesaran, M. H. (2015). Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors. *Journal of Econometrics*, 188(2), 393-420.

Diewert, W. E., & Fox, K. J. (2022). Substitution Bias in Multilateral Methods for CPI Construction. *Journal of Business & Economic Statistics*, 40(1), 296-311.

Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4), 549-560.

Gleason, K., Zhang, Y., & Wilson, A. (2026). Platform Design and Price Parity: The Empirical Impact of Buy Box Suppression. *Journal of Industrial Economics*.

Hausman, J. A. (1978). Specification Tests in Econometrics. *Econometrica*, 46(6), 1251-1271.

InvasWMS (2024). Costos Logísticos en Latinoamérica: Redúcelos con Tecnología. *Retrieved from invaswms.com*

Jararweh, Y., et al. (2025). Dimensional Weight Pricing and its impact on the e-commerce supply chain. *Transportation Research Part E: Logistics and Transportation Review*.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1), 161-182.

Jordà, Ò. (2023). Local Projections for Applied Economics. *Annual Review of Economics*, 15, 607-631.

Juodis, A., Karabiyik, H., & Westerlund, J. (2022). On the robustness of the pooled CCE estimator. *Journal of Econometrics*, 227(2), 241-267.

Kanuri, V. K., et al. (2025). Disentangling the Customer-Level, Cross-Channel Effects of Large-Order-Advantaged Online Shipping Policies. *MIS Quarterly*.

Lee, K., & Musolff, L. (2025). Quantifying Platform Penalties: The Value of Visibility in Digital Marketplaces. *Management Science*.

Li, D., Plagborg-Møller, M., & Wolf, C. K. (2024). Local projections vs. VARs: Lessons from thousands of DGPs. *Journal of Econometrics*, 244(2), 105658.

Locus (2024). Last-Mile Delivery in 2026: Costs, Challenges, Fixes. *Retrieved from locus.sh*

Maddala, G. S., & Wu, S. (1999). A Comparative Study of Unit Root Tests with Panel Data and a New Simple Test. *Oxford Bulletin of Economics and Statistics*, 61(S1), 631-652.

MegaMinds (2024). Why Last-Mile Delivery Is Still the Biggest Challenge in Logistics. *Retrieved from megamindstechnologies.com*

Montiel Olea, J. L., & Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4), 1789-1823.

Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4), 967-1012.

Pesaran, M. H. (2021). General diagnostic tests for cross-sectional dependence in panels. *Empirical Economics*, 60(1), 13-50.

Roy, R., Togawa, T., Park, J., Ishii, K., & Naidoo, V. (2024). The haptic sensation of heaviness conditions pricing preferences: Partitioned vs. combined pricing. *Journal of Business Research*, 173.

Singh, A., & Ardjmand, E. (2021). Dimensional weight pricing and its effect on packaging and logistics. *Journal of Business Logistics*, 42(3), 321-340.

Thakurele, A., Lu, S., & Xiao, P. (2024). The Impact of Platform Algorithm on Seller Pricing and Prominence. *Marketing Science*, Forthcoming.

Tran, T. (2024). Consumer Welfare and Partitioned Pricing in Digital Markets. *Journal of Consumer Psychology.*

Wang, Z., Nan, G., Kou, G., & Li, M. (2023). The impact of shipping fees on channel integration: A game-theoretic analysis. *European Journal of Operational Research*, 307(1), 162-177.

Appendix A: Tables

Table 2: Dynamic Estimates of the Logistics Impact on the Base Price (Light-Goods Cluster)

	Horizonte	Efecto_Precio	IC_Inferior	IC_Superior
1	0	-0.0006	-0.0022	0.0010
2	1	-0.0004	-0.0022	0.0014
3	2	-0.0002	-0.0022	0.0018
4	3	-0.0002	-0.0023	0.0019
5	4	-0.0008	-0.0031	0.0015
6	5	-0.0010	-0.0035	0.0014
7	6	-0.0010	-0.0037	0.0016
8	7	-0.0004	-0.0032	0.0024
9	8	-0.0003	-0.0033	0.0026
10	9	-0.0013	-0.0044	0.0017
11	10	-0.0020	-0.0051	0.0012
12	11	-0.0024	-0.0057	0.0010
13	12	-0.0023	-0.0058	0.0011
14	13	-0.0028	-0.0064	0.0007
15	14	-0.0033	-0.0069	0.0003

Note: The table reports the impact coefficients for the logarithm of the base price ($\ln P_t$) over a projected horizon of 0 to 14 days. The confidence intervals ($IC\_Inferior$ and $IC\_Superior$) are estimated at the 90% confidence level using Newey-West HAC robust standard errors and controlling for Two-Way Fixed Effects (TWFE).

Table 3: Dynamic Estimates of the Logistics Impact on Demand (Light-Goods Cluster)

	Horizonte	Efecto_Demanda	IC_Inferior	IC_Superior
1	0	-0.3605	-0.4024	-0.3186
2	1	-0.0961	-0.1328	-0.0593
3	2	-0.0890	-0.1253	-0.0526
4	3	-0.0624	-0.1013	-0.0235
5	4	-0.0703	-0.1067	-0.0339
6	5	-0.0871	-0.1262	-0.0479
7	6	-0.0710	-0.1084	-0.0336
8	7	-0.0941	-0.1347	-0.0534
9	8	-0.0587	-0.0978	-0.0196
10	9	-0.0377	-0.0757	0.0002
11	10	-0.0536	-0.0933	-0.0139
12	11	-0.0615	-0.0999	-0.0232
13	12	-0.0539	-0.0949	-0.0129
14	13	-0.0576	-0.1023	-0.0130
15	14	-0.0431	-0.0883	0.0021

Note: The table reports the impact coefficients for transformed quantity demanded ($IHS(Q_t)$) over a projected horizon of 0 to 14 days. Confidence intervals are estimated at the 90% confidence level using Newey-West HAC robust standard errors.

Table 4: Dynamic Estimates of the Logistics Impact on the Base Price (Heavy and Bulky Goods Cluster)

	Horizonte	Efecto_Precio	IC_Inferior	IC_Superior
1	0	-0.0004	-0.0019	0.0012
2	1	-0.0005	-0.0020	0.0010
3	2	-0.0013	-0.0030	0.0003
4	3	-0.0012	-0.0029	0.0005
5	4	-0.0008	-0.0026	0.0010
6	5	-0.0001	-0.0019	0.0017
7	6	0.0007	-0.0013	0.0027
8	7	0.0010	-0.0011	0.0031
9	8	0.0009	-0.0012	0.0030
10	9	0.0011	-0.0011	0.0033
11	10	0.0011	-0.0012	0.0033
12	11	0.0014	-0.0009	0.0037
13	12	0.0013	-0.0012	0.0037
14	13	0.0008	-0.0016	0.0033
15	14	0.0005	-0.0019	0.0030

Note: The table reports the impact coefficients for the logarithm of the base price ($\ln P_t$) over a projected horizon of 0 to 14 days. The confidence intervals (*IC_ Inferior* and *IC_ Superior*) are estimated at the 90% confidence level (TWFE + Newey-West).

Table 5: Dynamic Estimates of the Logistics Impact on Demand (Heavy and Bulky Goods Cluster)

	Horizonte	Efecto_Demanda	IC_Inferior	IC_Superior
1	0	-0.3303	-0.3793	-0.2812
2	1	-0.1331	-0.1708	-0.0955
3	2	-0.0646	-0.0999	-0.0294
4	3	-0.1008	-0.1385	-0.0631
5	4	-0.0787	-0.1146	-0.0428
6	5	-0.0808	-0.1158	-0.0457
7	6	-0.0856	-0.1234	-0.0477
8	7	-0.0675	-0.1081	-0.0268
9	8	-0.0569	-0.0965	-0.0173
10	9	-0.0676	-0.1088	-0.0263
11	10	-0.0407	-0.0782	-0.0033
12	11	-0.0410	-0.0754	-0.0065
13	12	-0.0492	-0.0851	-0.0134
14	13	-0.0722	-0.1128	-0.0315
15	14	-0.0511	-0.0937	-0.0085

Note: The table reports the impact coefficients for transformed quantity demanded ($IHS(Q_t)$) over a projected horizon of 0 to 14 days. Confidence intervals are estimated at the 90% confidence level using Newey-West HAC robust standard errors.

Declaración sobre el uso de Inteligencia Artificial

En cumplimiento con los principios de transparencia e integridad académica, el autor declara que durante el desarrollo de este Trabajo de Fin de Máster se han empleado herramientas basadas en modelos de lenguaje de Inteligencia Artificial (IA) como instrumentos de apoyo.

El uso de estas tecnologías se ha limitado de manera estricta a labores de asistencia metodológica y optimización del flujo de trabajo, concretamente en las siguientes áreas:

- **Asistencia en la redacción y traducción:** Apoyo en la mejora de la fluidez estilística, corrección ortotipográfica y traducción de fragmentos específicos para garantizar la claridad expositiva del texto.
- **Búsqueda y contraste de literatura:** Asistencia en la identificación preliminar de fuentes bibliográficas, mapeo de literatura relevante y contraste de conceptos teóricos.
- **Revisión y depuración de código:** Soporte en la identificación de errores de sintaxis, optimización de rutinas y revisión del código de programación empleado para el análisis empírico.

Se certifica que la concepción de la pregunta de investigación, el diseño del marco metodológico, el análisis crítico de la literatura, la interpretación de los resultados empíricos y las conclusiones derivadas de este estudio son obra original y exclusiva del autor. El autor ha revisado minuciosamente todos los resultados sugeridos por las herramientas de IA y asume la plena responsabilidad sobre la originalidad, veracidad y el contenido final del presente documento.