

The Price of Delay: Supply Chain Disruptions and Pricing Dynamics

Salomé Baslandze and Simon Fuchs

Working Paper 2025-8a

September 2025 (Revised March 2026)

Abstract: We study how supply chain disruptions shape consumer prices. We develop a simple model showing how disruptions that cause delivery delays affect pricing through product availability, beyond standard cost pass-through, and extend it to a setting with strategic interactions across firms. Empirically, we construct a novel micro-level dataset linking shipment-level US import records from bills of lading to granular consumer prices from the Numerator panel. Using delivery shortfalls, port congestion, and import costs in a shift-share identification strategy, we estimate the pass-through of supply chain shocks to retail prices. We find sizable but incomplete pass-through: both imported-input cost shocks and delivery delays raise prices, with stronger responses when disruptions persist. Firms also raise prices when competitors face disruptions—including those that do not import directly—indicating strategic pricing spillovers. Combining these estimates with back-of-the-envelope calculations from the model, we show that such interactions substantially amplify the aggregate price effects of supply chain shocks, particularly those operating through product availability, during the pandemic.

JEL classification: E31, F14

Key words: supply chains, inflation, delivery delays, strategic interactions, pass-through, inventory

<https://doi.org/10.29338/wp2025-08>

The authors are grateful to Valerie Nguyen, KC Pringle, and Michael Sparks for outstanding research assistance. They thank seminar and conference participants at the 2026 AEA/ASSA meetings, the University of Mississippi, Monash University, University of Melbourne, Reserve Bank of Australia, ISOT (World Bank), the SEM Conference (Georgia Tech), the Firm Dynamics and the Macroeconomy conference (Federal Reserve Bank of Atlanta), the Virtual Seminar Series in Innovation and Growth, the Federal Reserve Bank of Minneapolis, SED 2025, the BSE Summer Forum, Workshop on International Economic Networks (WIEN), Armenian Economic Association, and the System Committee on Macroeconomics, and 4th Conference on “Trade, value chains and financial linkages in the global economy” for helpful feedback. They are especially indebted to our discussants, Vanessa Alviarez, Davin Chor, Sebastian Heise, and Viacheslav Sheremirov, and to Leo Feler for support with the Numerator data. We also benefited from insightful comments by Pol Antràs, Michael McMahon, and Jon Willis. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Atlanta, the Federal Reserve System, or its staff.

Please address questions regarding content to Salomé Baslandze, Federal Reserve Bank of Atlanta and CEPR, baslandze.salome@gmail.com; or Simon Fuchs, Federal Reserve Bank of Atlanta, sfuchs.de@gmail.com.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed’s website at www.frbatlanta.org. Click “Publications” and then “Working Papers.” To receive e-mail notifications about new papers, use atlantafed.org/forms/subscribe.

1 Introduction


Throughout the pandemic, supply chain disruptions and their impact on inflation have been at the forefront of policy discussions. International trade and supply chains were disrupted in many different ways, ranging from backlogs at ports, heightened maritime shipping costs, production disruptions at the origin due to COVID restrictions, as well as increased demand shocks that put strains on limited production capacity. These various measures of disruptions are summarized by the NY Fed Global Supply Chain Pressure Index (GSCPI) depicted in blue in Figure 1. At the same time, US inflation has experienced historically high levels, as illustrated by CPI inflation depicted in Figure 1 in red. While a growing literature links supply chain disruptions to observed inflation,² micro-level evidence remains limited. In particular, it is unclear how strongly firms pass these disruptions through to *consumer prices* and what role *strategic interactions* play in shaping the resulting price dynamics.

From a policy perspective, the 2021–2022 inflation surge raised a central question: to what extent did rising prices reflect demand expansion versus supply constraints? Because standard stabilization tools operate primarily on demand, both the appropriate response and the expected persistence of inflation depend on this distinction. Our contribution speaks to the supply side of this debate. We provide micro-level evidence on two supply-chain channels through which disruptions affect consumer prices: (i) imported-input and shipping cost shocks and (ii) delivery-related tightness affecting product availability and markups. We estimate the pass-through of firms’ own disruptions and examine how competitors’ disruptions amplify price responses through strategic interactions.

To study these mechanisms, we proceed in four steps. First, we develop a motivating model of delivery delays, product availability, and firm pricing to illustrate—in a transparent way—how supply chain pressures translate into higher prices and to provide the estimating framework. Second, we assemble a large-scale micro-level dataset combining firm–product-level prices with firm-level measures of delivery shortfalls, port congestion, and cost shocks from unit import costs and freight costs. Using these data, we then estimate how firms’ own supply chain disruptions—delivery shortfalls, unit import cost shocks, and freight cost shocks—affect pricing using OLS and a shift-share instrumental variable approach. Next, we extend the framework to incorporate strategic interactions and estimate the effect of competitors’ disruptions on firm pricing, over and above the firm’s own shocks. Finally, we use the estimated objects in a simple accounting exercise to show how supply chain disruptions could shape aggregate price dynamics through direct cost effects, own availability effects, rivals’ availability effects, and the amplification generated by strategic feedback.

²See, among others, [Guido Ascari, Dennis Bonam and Andra Smadu \(2024\)](#); [Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li and Francesco Zanetti \(2024\)](#); [David Finck, Mathias Klein and Peter Tillmann \(2024\)](#); [Giulia Brancaccio, Myrto Kalouptsi and Theodore Papageorgiou \(2025\)](#).

Figure 1. Supply Chain Pressure Index and Inflation



Notes: Panel (A) plots the NY Fed Global Supply Chain Pressure Index (GSCPI). Panel (B) plots U.S. CPI inflation for all items less food and energy (12-month percent change; BLS/FRED series CPILFESL). The figure illustrates the close co-movement between aggregate supply-chain pressure and core inflation over the pandemic period.

We begin by developing a simple motivating model to illustrate how delivery delays that affect product availability shape firm pricing beyond standard cost pass-through. The model delivers an optimal inventory–pricing benchmark that links replenishment lead times and base-stock targets to product availability and markups. Product availability—the probability that an arriving customer is served—falls with longer lead times or lower base-stock and rises with faster replenishment. We show that lower availability reduces the effective price elasticity faced by the firm, which raises the optimal markup, while higher availability increases the effective elasticity and compresses markups. Consequently, supply chain shocks through longer delays and lower availability lead to higher prices.

We use our simple model to deliver reduced-form expressions for firms’ optimal price dynamics that split observed price changes into (i) movements in marginal cost and (ii) movements in product availability. We map import delivery shortfalls to changes in availability and decompose import-based cost shifters into unit import and freight cost components. The model also provides qualitative predictions: higher delivery shortfalls lead to lower product availability, lowering the effective price elasticity and increasing markups, while increases in cost push prices up with a pass-through strictly between zero and one.

To trace how upstream supply chain disruptions affect downstream consumer prices, we build a matched firm–product panel that combines two unusually granular sources: shipment-level import records, which reveal a firm’s supply pipeline at high frequency, and receipt-level consumer transactions, which record the prices actually paid by consumers for specific products over time.

On the upstream side, we use S&P Global Panjiva Bills of Lading (2007–2023), a transaction-level dataset of U.S. maritime imports reporting shipper–consignee pairs, HS descriptions, quantities/weights,

and vessel arrival dates.³ Event-level time stamps and counterparty identifiers allow us to construct firm $\times$ imported-product-code (HS) measures of delivery shortfalls—defined as (negative of) deviations in k -month cumulative import volumes relative to the 2019 pre-pandemic benchmark, which we then aggregate to the firm level. Hence, large positive shortfalls indicate larger downward deviations in firms’ imports relative to the pre-pandemic baseline. We complement our measures of delivery shortfalls with firm exposure to port congestion (dwell times) by combining data on a firm’s port usage with U.S. port dwell time data from [Fuchs and Wong \(2022\)](#).⁴ Finally, we assemble two model-motivated marginal-cost shifters: unit import costs and freight costs. Changes in import costs (unit values at the firm $\times$ HS-code level) and freight costs (at the port level) are aggregated using the firm’s predetermined exposures to HS codes and ports to derive firm-specific shift-share shocks.

To measure downstream consumer prices, we use Numerator’s Consumer Panel (2019–2023), which records itemized receipts from both brick-and-mortar retailers and online. The data span a wide range of goods in the consumer basket, covering many product categories with detailed information on quantities, sales, producer names, brands, and purchase timestamps. We define products at the firm–brand–product category level and track 12-month price changes as our baseline outcome. The receipt-level detail ensures we capture realized transaction prices, including promotions and coupons, across households, channels, and geographies. Aggregate spending in the panel closely tracks Census retail sales, and the implied inflation measures comove tightly with CPI inflation and with inflation from a comparable basket of goods, supporting the reliability of the data for studying consumer price dynamics.

We link downstream consumer prices to upstream import activity by matching firms in the Numerator consumer panel to their international shipments in the Panjiva data. After extensive name cleaning and standardization across datasets, this procedure yields a novel firm–product dataset that connects U.S. consumer prices to firms’ import exposures. The matched sample identifies consumer-product firms that import in Panjiva,⁵ accounting for more than 40% of product–month observations in Numerator and about half of total sales.

Two descriptive patterns emerge. First, delivery shortfalls during 2020–2022 were large but highly uneven—both across product categories and across firms within a category. Second, price changes display similarly wide dispersion: larger firms raise prices more on average, while more diversified

³Panjiva is based on the near-universe of waterborne bills of lading filed with U.S. Customs and Border Protection, providing named shipper–consignee identifiers and event-time stamps for individual shipments.

⁴Similar data have also recently been used by [Giulia Brancaccio, Myrto Kalouptsidi and Theodore Papageorgiou \(2024, 2025\)](#); [Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li and Francesco Zanetti \(2024\)](#) to study supply chain disruptions, but to the best of our knowledge this is the first study that explicitly constructs and exploits heterogeneous incidence at the firm-level.

⁵As described in more detail in the data section, we focus on a stable sample of importing firms to reduce data noise.

importers experience more muted price growth.

Using our matched firm–product panel linking shipment-level imports to receipt-level consumer prices, we estimate how supply-chain disruptions pass through to retail pricing. Guided by the model, we distinguish three channels through which disruptions affect prices. First, a *marginal-cost* channel reflecting firms’ exposure to changes in *own* imported-input prices and transportation costs, proxied by changes in HS-level unit import values and port-level freight and insurance costs, and aggregated to the firm using predetermined (2019) HS import and port-exposure weights. Second, the *own availability* channel resulting from disruptions that reduce the probability that products are in stock—either because imports of final goods decline or because shortages of imported inputs constrain production—proxied by firm-level delivery shortfalls. Third, *within-market spillovers*, reflecting strategic interactions across firms, captured by market-share-weighted competitor indices that measure exposure to *rivals’ cost and availability shocks*.

Starting with the first two channels—the pass-through of firms’ *own* disruptions into prices—we estimate a baseline OLS specification that regresses 12-month product-level price changes on firm-level delivery shortfalls and import-based marginal-cost shifters. On their own, delivery shortfalls reflect not only supply constraints but also aggregate and idiosyncratic demand factors: stronger aggregate or firm-level demand increases imports and therefore mechanically lowers measured shortfalls. To account for common market-wide demand (and supply) movements, we include product-category×time fixed effects. We also control for firm fixed effects to absorb time-invariant differences in pricing across firms. This specification therefore identifies pass-through coefficients by comparing price changes across firms within the same market and time that face different delivery shortfalls and cost exposures.

However, idiosyncratic demand shocks may still bias the OLS estimates of the shortfall coefficient. For example, a positive firm-specific demand shock may both prompt a firm to accelerate orders—mechanically reducing measured shortfalls—and raise prices, thereby biasing the estimated pass-through coefficient toward zero.⁶ To address this concern, we implement a shift-share IV strategy that instruments firm-level delivery shortfalls using predicted exposure measures based on firms’ pre-pandemic (2019) import portfolios and transportation patterns. Specifically, we combine firms’ fixed shares across HS codes and importing ports with common supply shocks: HS-level disruption shifters constructed from leave-one-out shortfalls and port-level congestion shocks measured by changes in dwell times. Conditional on firm fixed effects and product-category×quarter fixed effects, identification comes from comparing firms in the same product market and time period but with different pre-pandemic sourcing patterns across products and ports. The instruments therefore isolate variation in delivery shortfalls driven by pre-

⁶Measurement error in the shortfall proxy can generate similar attenuation.

determined exposure to common supply disruptions rather than firm-specific demand conditions. In practice, the instruments are strong, with first-stage F-statistics well above conventional thresholds.

We find that both delivery shortfalls and imported-input cost shocks pass through to consumer prices in sizable but incomplete manner. Unit import costs and freight costs exhibit positive pass-through of about 0.13–0.15, rising to 0.19–0.23 when shocks are lagged by one month, indicating that cost shocks transmit with a short delay.⁷ Delivery shortfalls (three-month moving averages) show similarly sizable effects. As expected, OLS estimates understate their impact, but IV specifications yield stable elasticities around 0.20. Pass-through also strengthens as disruptions persist across longer horizons, suggesting firms adjust prices more when shocks appear durable. We further explore heterogeneity across product types and firms. Although durable goods experienced stronger price growth during the pandemic, this largely reflects demand pressures rather than higher supply-side pass-through. Multi-national firms—defined as those with related-party import transactions in Panjiva—show no systematic differences in pass-through relative to other firms. Overall, the results indicate that both cost shocks and availability constraints contributed meaningfully to consumer price inflation, with the IV strategy ensuring the estimates reflect supply-driven rather than demand-driven effects in delivery shortfalls.

We next examine strategic interactions to assess whether firms’ price responses to supply chain shocks were amplified by their rivals’ disruptions. Extending the simple model to allow for strategic interactions (Amiti, Itskhoki and Konings, 2019) highlights how each firm sets prices not only based on its own costs and availability but also on competitors’ conditions. Empirically, we augment our pass-through specification to include rivals’ market-share–weighted delivery shortfalls and cost shocks (unit and freight), defining rivals within finely-delineated consumer-product categories.

The results show that, beyond the direct effects of firms’ own disruptions, supply shocks also transmit through strategic channels. Firms raise prices in response to competitors’ disruptions, with the pass-through of rivals’ shortfalls to own prices about half the size of the own-pass-through elasticity. Competitors’ cost shocks likewise spill over, indicating that supply bottlenecks and cost pressures propagate through competitive interactions. These effects extend even to *non-importers*, who raise prices when importing rivals are disrupted, consistent with demand shifting toward unaffected sellers and weaker competitive discipline. Finally, both own and competitor pass-through are state-dependent: price responses are significantly larger in high-inflation sectors and during periods of elevated aggregate inflation. Together, these findings imply that supply chain disruptions can generate aggregate price pressures well beyond directly affected firms, with amplification that is strongest when inflation is already elevated.

⁷For comparability with estimates in the cost pass-through literature, our import-cost exposures are scaled by industry import-intensity weights, and shortfalls are scaled by intermediates import shares.

Finally, we conduct a model-consistent accounting exercise in an impulse–response spirit, feeding the monthly paths of our supply-chain shifters into the estimated micro pricing system to trace how they propagate through market shares and strategic interactions into consumer prices. The results show that while both cost and availability shocks raise prices, availability disruptions have the largest effects. Because these disruptions are widespread, strategic interactions amplify them and propagate price increases across firms, including those with little direct exposure. Interpreted as an impulse–response exercise—holding exposures, pass-through parameters, and market shares fixed—these findings highlight the central role of availability constraints and the amplification that arises when many firms face the same bottlenecks.

Literature. Our paper contributes to several areas of the literature. First, we speak to the pass-through literature for internationally sourcing firms by placing availability via delivery shortfalls—a reduced-form proxy for delivery delays and inventory constraints—on the same pricing margin as marginal costs. Cost-focused work shows that tariff cuts reduce marginal costs yet raise markups, implying incomplete pass-through ([De Loecker et al., 2016](#)); related evidence on energy shocks documents full pass-through for cost increases but only partial pass-through for cost reductions, with limited aggregate inflation effects given energy’s small cost share and substantial within-industry heterogeneity ([Lafrogne-Joussier, Martin and Mejean, 2023](#)). On the strategic side, [Amiti, Itskhoki and Konings \(2019\)](#) document strong strategic complementarities in price setting, and [Albagli et al. \(2025\)](#) show these complementarities can dominate cost forces and are state-dependent; [Lu et al. \(2024\)](#) provide related evidence on how market power and nominal rigidities shape cost pass-through and inflation dynamics. We fold these insights into a unified specification that treats own costs, competitors’ prices, and—critically—own and rivals’ availability in parallel, thus isolating a distinct mechanism whereby supply-chain disruptions via delivery delays shift prices in addition to costs. Our contribution is to provide novel evidence that availability constraints arising from supply-chain disruptions have significant effects on firm pricing—not only through firms’ own availability shocks, but also through strategic complementarities with rivals’ disruptions—over and above standard cost pass-through effects.

We also contribute to the literature that links inventories, delivery frictions, and pricing. Dynamic inventory models with monopolistic competition show that scarce stocks raise markups and slow pass-through, with delivery lags and fixed order costs generating state-dependent pricing and precautionary inventory behavior ([Alessandria, Kaboski and Midrigan, 2010](#); [Alessandria et al., 2023](#)). Work on sourcing under stochastic delivery times emphasizes that binding inventory constraints tilt pricing toward precautionary markups and induce higher buffers when supply is slower or riskier ([Carreras-Valle and Ferrari, 2025](#)). Macro and retail studies with inventories and nominal frictions likewise pro-

duce real rigidities, inventory-driven sales, and state-dependent markups ([Khan and Thomas, 2007](#); [Kryvtsov and Midrigan, 2013](#); [Aguirregabiria, 1999](#)). Relative to these contributions, our approach is deliberately minimal and analytically transparent: it delivers closed-form objects that pin down an availability-adjusted markup and a sharp, estimable decomposition of price changes into cost and availability channels. This tractability directly motivates reduced-form empirical specifications for both own pass-through and strategic exposure to rivals' conditions, and it provides a clean bridge to frameworks that measure strategic complementarities in price setting—most notably extending [Amiti, Itskhoki and Konings \(2019\)](#) by bringing firm and rival availability into the same empirical mapping as costs and competitor prices.

Finally, we contribute to a string of recent research that links pandemic-era bottlenecks to inflation and macro dynamics. Survey evidence shows disruptions shifting expected unit costs ([Meyer, Prescott and Sheng, 2023](#)); New Keynesian analyses with binding capacity constraints trace upward shifts in Phillips curves and amplified goods-sector inflation ([Comin, Johnson and Jones, 2023](#)); indices based on container-ship movements and DSGE exercises attribute a sizable 2021 inflation impulse to supply-chain shocks and study monetary-policy interactions ([Bai et al., 2024](#); [Amiti et al., 2024](#)); and high-frequency micro evidence documents how pandemic-induced stockouts transmitted into temporary but significant inflationary pressures ([Cavallo and Kryvtsov, 2023](#)). Related work emphasizes the central role of transportation rigidities—ships and ports—in propagating shocks to trade costs and inflation ([Brancaccio, Kalouptsi and Papageorgiou, 2025](#)). At the micro level, [Liu, Smirnyagin and Tsyvinski \(2024\)](#) show that delivery lapses depress supplier balance sheets and firm performance without tracing consumer-price pass-through, while [Borusyak and Jaravel \(2021\)](#) study tariff pass-through but not pandemic-related delays. Our contribution is to provide causal evidence—using unique micro-level data that link firm-specific delivery shortfalls to product-level prices—on how supply-chain disruptions move consumer prices and how strategic complementarities across firms amplify these effects. We isolate delay-specific shocks from input-cost movements and pair the evidence with a tractable model-to-data mapping that delivers closed-form pass-through objects and a transparent decomposition of price changes into cost and availability channels, quantifying both own pass-through and strategic exposure to rivals' conditions.

The remainder of the paper is structured as follows. Section 2 introduces a motivating model. Section 3 describes the data construction and provides summary statistics. Section 4 analyzes the pass-through of both own and competitors' supply chain disruptions to consumer prices. Section 5 presents a simple accounting exercise to quantify potential implications of our pass-through estimates. Section 6 concludes.

2 Theory: A Model of Pricing with Supply Chain Disruptions

In this section, we present a simple and transparent model linking supply-chain conditions to firms' pricing and inventory choices. While the pass-through of cost shocks to prices is well understood, the model highlights a distinct channel through which disruptions to (import) deliveries affect prices by altering product availability. A downstream firm facing flow demand and stochastic replenishment chooses a posted price and a base-stock target. Under a base-stock policy, we obtain closed-form expressions for the stock-in probability and a generalized Lerner rule in which an availability wedge (driven by congestion in replenishment) raises markups when stock is scarce. Linearizing the pricing condition delivers a two-term decomposition of price changes into a marginal-cost component and an availability component—a structure we take directly to the data. The derivations and notation are collected in Appendix A.1; in Section 4.2 we extend the model to allow for strategic interactions between firms.

2.1 Inventory and Pricing under Delivery Frictions

We study a single downstream firm that *procures* a ready-made product from an upstream supplier rather than producing it in-house.⁸ The firm chooses a posted (log) price p and a base-stock target τ , our *availability shifter* that summarizes order size, review frequency, expediting, and sourcing intensity. Operationally, a base-stock (order-up-to) rule aims to restore on-hand inventory to τ whenever a delivery arrives; higher τ raises the chance that an arriving customer is served but increases holding cost, so τ is the reduced-form buffer-capacity lever given the upstream replenishment process.⁹

Customer demand arrives as a flow at rate $\lambda(p)$, decreasing in price. Replenishment of the firm's inventory occurs through the supplier with *stochastic* delivery times summarized by a *replenishment speed* $\mu > 0$: higher μ means faster deliveries (mean lead time $1/\mu$). Only demand arrivals that find stock on hand are served. Let p and c denote *log* price and *log* marginal cost, with corresponding *levels* $P := e^p$ and $MC := e^c$, and let $h > 0$ be the holding-cost rate per unit of average on-hand inventory $\mathbb{E}[I(\tau, \lambda(p), \mu)]$. For tractability, we treat the imported procurement cost as constant per unit in the benchmark, so the imported component's marginal and average procurement costs coincide. Finally, we assume that customers who arrive when no stock is on hand are not served, an assumption commonly

⁸Although we maintain this interpretation, the model can readily be adapted to view firms as receiving intermediate goods deliveries; both interpretations map naturally to our import data discussed later.

⁹We deliberately abstract from several operational features to keep the mapping transparent: a single product under continuous review with lost sales (rather than backlogging), no fixed order or setup costs or minimum order quantities (which would induce (s, S) policies), and no storage limits, perishability, or order crossover. We also treat the replenishment speed μ as within-period exogenous. As we will show below, these simplifications deliver the closed-form acceptance probability $s(\tau, \lambda, \mu)$ and a clean link between pricing and inventory that we can take to the data. Formal details and variants are provided in Appendix A.1.

called the *lost-sales setting*.¹⁰ Focusing on steady-state flow payoffs in a lost-sales setting, the firm's static profit rate is¹¹

$$\Pi(p, \tau) = (P - MC) \lambda(p) s(\tau, \lambda(p), \mu) - h \mathbb{E}[I(\tau, \lambda(p), \mu)],$$

where $s(\tau, \lambda, \mu) \in [0, 1]$ is the acceptance (stock-in) fraction—the probability that a demand arrival is served.

To obtain closed-form objects that map cleanly to the data, we adopt two standard assumptions. (i) Flow demand is log-linear, $\lambda(p) = \Lambda e^{-\sigma p}$, so $\sigma > 0$ is the demand semi-elasticity, and $\Lambda > 0$ collects demand shifters. (ii) Replenishment lead times are exponential with rate $\mu > 0$ (mean $1/\mu$), and the firm follows a base-stock rule at level $\tau \in \mathbb{N}$.¹² Define

$$r \equiv \frac{\lambda}{\lambda + \mu} \in (0, 1), \quad s(\tau, \lambda, \mu) = 1 - r^\tau,$$

so r measures how fast demand arrives relative to replenishment (“congestion”), and one can show that under the assumptions above s is the probability an arriving customer finds stock on hand - our closed-form stock-in probability.¹³ It is also convenient to track how demand pressure erodes availability via

$$\kappa(\tau, \lambda, \mu) \equiv \lambda \frac{\partial s}{\partial \lambda} = -\tau(1 - r)r^\tau < 0,$$

an *availability adjustment* that will enter the pricing condition. Intuitively, r summarizes congestion, s summarizes how often customers are served, and κ summarizes how sensitive that service probability is to demand pressure—three workhorse objects that link lead times and inventory choices to prices and that carry through to the empirical mapping below.

Under these stochastic assumptions, we have derived a mildly adjusted standard firm optimization

¹⁰The model's baseline mapping treats stockouts as lost sales. The assumption is made for analytical convenience. In some markets—especially for more durable goods—unmet demand may instead be deferred (e.g., backorders or delayed purchases), making pricing partly dynamic through a backlog state. In such settings, availability shocks can affect prices through both contemporaneous scarcity and intertemporal demand reallocation, potentially changing persistence (and in some cases amplifying price responses if customers wait rather than substitute). Appendix A.3.3 provides an extension of the baseline model under partially deferred consumption where the availability wedge is attenuated by the strength of deferral and vanishes in the limit where sales can be fully deferred.

¹¹Lemma 1 in Appendix A.1 shows that the profit rate is an equivalent objective to long-run average profit under the memoryless property of the exponential distribution, (A23)–(A25).

¹²A realistic microfoundation for exponential lead times comes from the assumptions that upstream suppliers fill orders subject to capacity constraints, with orders that are waiting to be served forming a queue. Formally, we can be generic about both the arrival and servicing process of that upstream queue as long as we accept an asymptotic characterization of the queue under heavy traffic. In queuing notation, we consider a upstream G/G/1 (single-server queue with interarrival and servicing times having an arbitrary distribution) queue in heavy traffic: the stationary lead time is approximately $\text{Exp}(\mu_{\text{HT}})$ with μ_{HT} proportional to system slack and dampened by variability; see Appendix A.1, Lemma 2.

¹³Appendix A.1 first derives the stock-in probability from the exponential arrival processes for customers and upstream suppliers and then furthermore shows that the arrival probability for any individual customers, s , equals the time-average availability use in the objective function, the profit rate (Lemma 1).

problem. The firm chooses a profit-maximizing price and, in our setting, also chooses a base stock to maintain product availability and internalizes the effect of pricing on availability. Given this setup, we are now in a position to characterize the firm's optimal joint pricing and inventory choice. The following proposition summarizes the firm's optimal choice under the additional assumption that pricing is constant within a replenishment cycle.¹⁴

Proposition 1 (Optimal pricing and inventory in the simple model). *For any τ , the profit-maximizing price satisfies the generalized Lerner condition in levels:*

$$\frac{P^* - MC}{P^*} = \frac{1}{\sigma \left(1 + \frac{\kappa}{s}\right)} \quad \left(\text{equivalently } p^* = c - \ln \left[1 - \frac{1}{\sigma(1 + \kappa/s)}\right]\right),$$

with $s = 1 - r^\tau$, $\kappa = -\tau(1 - r)r^\tau$, $\lambda = \Lambda e^{-\sigma P^*}$, and $r = \lambda/(\lambda + \mu)$. If τ is chosen discretely, an interior base-stock τ^* solves the exact marginal condition

$$(P^* - MC) \lambda r^{\tau^*} (1 - r) = h(1 - r^{\tau^*+1}),$$

with $\lambda = \Lambda e^{-\sigma P^*}$. Under the continuous relaxation of τ , this first-order condition admits a unique solution with $r^{\tau^*} \in (0, 1)$.

These results have a simple interpretation. The pricing rule is a generalized Lerner condition in which the usual elasticity σ is adjusted by the *availability term* $1 + \kappa/s$. Because $\kappa/s < 0$ whenever stockouts can occur, the effective elasticity $\sigma(1 + \kappa/s)$ is below σ , implying a markup fraction $(P^* - MC)/P^*$ above the frictionless benchmark $1/\sigma$. The inventory condition trades off the marginal contribution from an additional unit of buffer against its holding cost: an extra unit pays off exactly when, before the next delivery, the next τ superposed events are $\tau - 1$ demand arrivals followed by a delivery, an event with probability $r^\tau(1 - r)$, yielding marginal revenue $(P^* - MC)\lambda r^\tau(1 - r)$ that is set equal to $h(1 - r^{\tau+1})$. As replenishment becomes faster, that event becomes less likely, the shadow value of inventory falls, and the optimal base-stock τ^* declines. Taken together, higher μ improves availability through two channels: it directly raises the acceptance probability s and, by pushing κ/s toward

¹⁴We assume prices are fixed within a replenishment cycle. This is consistent with micro evidence on infrequent price adjustment and with standard pricing frictions. For U.S. consumer prices, [Nakamura and Steinsson \(2008\)](#) show that once sales are netted out, the median frequency of nonsale price changes is on the order of 9–12% per month (implying median durations of roughly 8–11 months), whereas including sales yields much shorter durations. For the euro area, [Alvarez et al. \(2006\)](#) document average monthly frequencies near 15% (implying typical durations close to a year), again indicating that many prices remain unchanged for several months at a time. These facts make it empirically plausible that a firm does not reprice within a single replenishment cycle. The assumption also buys tractability: it delivers closed-form characterizations for availability, markups, and pass-through that map cleanly to reduced-form specifications.

zero, increases the effective elasticity and compresses the markup; both effects reduce the incentive to hold inventory.

Figure 2 provides a compact visualization of the two optimality margins in Proposition 1. Panel (a) illustrates the inventory path under delivery delays: inventory declines between replenishments and is restored upon shipment arrivals, so slower replenishment (lower μ , longer expected lead times) increases stockout risk and lowers the stock-in probability $s(\tau, \lambda, \mu)$. Panel (b) depicts the inventory-choice condition: the optimal base-stock τ equates the marginal revenue from one additional unit of buffer inventory, $MR^{inv}(\tau, \mu)$, to the marginal holding-cost schedule, $MC^{inv}(\tau)$; when delivery slows, $MR^{inv}(\tau, \mu)$ shifts up, and the optimal τ rises. Panel (c) shows the pricing condition in accepted units: the availability wedge shifts marginal revenue below the frictionless benchmark, so tighter availability (lower s) raises the markup and reduces the accepted flow; as availability improves ($s \rightarrow 1$), the outcome moves toward the frictionless point. Panel (d) summarizes the notation used in the figure.

Corollary 1 (Comparative statics with respect to lead times μ). *Holding primitives (Λ, σ, c, h) fixed,*


$$\partial_{\ln \mu} \ln s = \frac{\tau r^\tau (1-r)}{1-r^\tau} > 0, \quad \partial_{\ln \mu} p^*(\mu) < 0, \quad \partial_{\ln \mu} \tau^*(\mu) < 0.$$

In the fast-replenishment limit $\mu \rightarrow \infty$, we have $r \rightarrow 0$ and $s \rightarrow 1$; by the Lerner rule the markup fraction approaches $1/\sigma$, so the optimal level price tends to $MC \cdot \sigma/(\sigma - 1)$ and the optimal log price tends to $c + \ln(\sigma/(\sigma - 1))$. Conversely, in the slow-replenishment limit $\mu \rightarrow 0$, we have $1 - r \approx \mu/(\lambda + \mu)$ and hence $s \approx \tau \mu/(\lambda + \mu)$; in this region the price response becomes locally flat in μ with $\partial_{\ln \mu} p^ \rightarrow 0$.*

Corollary 1 says, in plain terms, that faster replenishment makes more arrivals get served, scarcity fades, and the firm faces a “softer” constraint—so markups and prices fall. At the same time, an additional buffer unit is less valuable (stockouts are less likely before the next delivery), so the optimal base-stock level declines. These forces bite most in the middle of the state space—when the system is neither almost always stocked nor almost always empty—because small improvements in delivery speed then move both availability and prices the most. At the extremes, effects are muted: with near-instant delivery, the frictionless markup prevails; with very slow delivery, availability is so low that further slowdowns barely move pricing. As shown in Figure 2, panel (b), faster delivery flattens the marginal revenue curve and lowers the optimal buffer.

To recap: Our simple model delivers two workhorse objects in closed form—a stock-in (availability) probability and an availability-adjusted markup rule—and uses them to generate intuitive comparative statics: faster replenishment raises availability, compresses markups, and lowers the value of holding inventory, while slower or riskier replenishment does the opposite. These closed-form objects also yield

Figure 2. Inventory Management with Delivery Delays



Notes: Panel (a) illustrates the inventory path under delivery delays. Inventory declines between replenishments and is restored upon shipment arrivals with expected inter-arrival time $1/\lambda$. A lower arrival rate increases stockout risk (blue region), reducing stock-in probability $s(\tau, \lambda)$. Panel (b) illustrates the optimal base-stock decision. The horizontal axis shows the base-stock (availability shifter) τ and the vertical axis shows marginal value and marginal holding cost per unit time. $MC^{inv}(\tau)$ denotes the marginal holding-cost schedule, while $MR^{inv}(\tau, \mu)$ denotes the marginal revenue from an additional unit of buffer inventory. When replenishment slows ($\mu^2 < \mu^1$), the marginal value of inventory rises and the optimal base-stock increases. Panel (c) illustrates pricing in accepted units. The horizontal axis shows accepted demand q^A and the vertical axis shows price p . Inverse demand is $p = D(q^A)$. The availability-adjusted marginal-revenue curve $MR^s(q^A)$ lies below the frictionless benchmark $MR^0(q^A)$ (when $s \equiv 1$). Lower availability (smaller s) reduces the effective elasticity and increases the markup, raising prices and lowering accepted quantities relative to the frictionless outcome.

a ready-to-estimate linear decomposition of price changes into cost and availability components, which underpins our empirical design. Relative to existing work on joint pricing and inventory, which typically solves dynamic models numerically, our approach keeps the same core mechanism but in a tractable, closed-form environment. For example, [Alessandria, Kaboski and Midrigan \(2010\)](#) study a dynamic importer with (S, s) ordering and delivery lags that link markups to the shadow value of inventories, while [George Alessandria, Shafaat Yar Khan, Armen Khederlarian, Carter Mix and Kim J. Ruhl \(2023\)](#) show in a heterogeneous-firm GE setting that uncertain shipping delays are contractionary and raise prices via

stockouts; our scarcity wedge captures the same channel in reduced form. Likewise, [Carreras-Valle and Ferrari \(2025\)](#) obtain precautionary markups under stochastic delivery times, mirroring our availability channel.¹⁵ Our contribution is to isolate these forces in a mildly extended standard firm optimization problem, delivering analytic expressions that map directly to reduced-form pass-through—both own and strategic—without solving a dynamic program. We next turn towards what the model implies for our empirical design.¹⁶

2.2 From Theory to Empirics

By Proposition 1, the optimal price is pinned down by two ingredients: marginal cost and an availability-adjusted scarcity term (the κ/s wedge) that compresses the effective demand elasticity when stockouts are likely. Linearizing around an operating point yields¹⁷

$$dp = \underbrace{\frac{1}{1+\Gamma}}_{:= \alpha \in (0,1)} dmc + \underbrace{\frac{\Lambda}{1+\Gamma}}_{:= \beta_s < 0} d \ln s + \varepsilon, \quad (\text{E1})$$

where α captures *own cost pass-through* and β_s captures *own availability pass-through*. Any price change feeds back into scarcity—raising p reduces arrivals, improves availability, and softens the desired markup, so *both* cost and availability shocks are attenuated by the same factor $\alpha \equiv 1/(1 + \Gamma)$. When this scarcity feedback is weak ($\Gamma \approx 0$), pass-through is close to one; when it is strong, pass-through is muted. Because $\beta_s = \alpha \Lambda$ with $\Lambda < 0$, better availability raises effective elasticity and lowers the desired markup, and the resulting price change is scaled by the same α . In short, costs shift the cost side, and availability shifts the markup side, but a single attenuation parameter α governs how quickly prices

¹⁵Furthermore, classic inventory macro models ([Khan and Thomas, 2007](#); [Kryvtsov and Midrigan, 2013](#)) imply similar comparative statics—faster replenishment compresses markups and weakens buffer motives—which we quantify in a tractable flow-demand, exponential-lead-time environment. See Appendix A.3 for an explicit comparison with [Kryvtsov and Midrigan \(2013\)](#).

¹⁶The model in this section is intentionally stylized (base-stock ordering, lost sales, and a posted price fixed within a replenishment cycle) to yield a transparent mapping from imported-input costs and availability into desired markups. Extensions that typically require numerical solutions—dynamic pricing and ordering, backlogging/deferred demand, non-exponential lead times, and multi-product capacity constraints—would mainly affect the timing and magnitude of pass-through (e.g., backlogging tends to dampen the scarcity/markup channel, while deferred demand or tighter capacity can amplify it) without changing intuition for the qualitative distinction between cost shifters and availability disruptions. See Appendix A for more details and discussion.

¹⁷Where we define the markup mapping $\mathcal{M}(p, \ln s) := \sigma(1 + z(p, \ln s))[\sigma(1 + z(p, \ln s))]^{-1}$ so that $\Gamma := -\partial \mathcal{M} / \partial p > 0$ and $\Lambda := \partial \mathcal{M} / \partial \ln s < 0$. Start from the Lerner rule in Proposition 1 and write the price–cost gap as a markup mapping

$$\mathcal{M}(p, \ln s) := -\ln\left(1 - \frac{1}{\sigma[1+z(p, \ln s)]}\right), \quad z := \kappa/s.$$

Then the optimality condition is $F(p, c, \ln s) := p - c - \mathcal{M}(p, \ln s) = 0$. Totally differentiating around an operating point gives $(1 + \Gamma)dp = dc + \Lambda d \ln s$, where $\Gamma := -\partial_p \mathcal{M}(p, \ln s) \geq 0$ and $\Lambda := \partial_{\ln s} \mathcal{M}(p, \ln s) \leq 0$. Rearranging yields the linear approximation $dp = \alpha dc + \beta_s d \ln s$ with $\alpha = 1/(1 + \Gamma)$ and $\beta_s = \Lambda/(1 + \Gamma)$. Replacing dc by dmc (the log marginal-cost change) gives (E1). Detailed derivations in Appendix A.1.

move.

From theory to measurement. To take the model to the data, we use two observables: (i) unit–import and freight price indices that move imported-input part of firm marginal cost, and (ii) a *delivery shortfall* index S_{it} that proxies for availability. Intuitively, shortfalls in imports slow replenishment, draw down inventories I , and reduce the stock–in probability s , which raises prices through the availability term in Proposition 1. Locally, we summarize this link by

$$\Delta \ln s_{it} \approx -\phi_i S_{it}, \quad \phi_i := \underbrace{\frac{\partial \ln s}{\partial \ln I}}_{\eta_{sI} > 0} \times \underbrace{\frac{\partial \ln I}{\partial \ln M}}_{\eta_{IM} \text{ (import dependence)}},$$

so the impact of a given shortfall is larger when availability is more sensitive to inventories and when inventories depend more on inbound flows.

On the cost side, we decompose marginal–cost changes as $\Delta mc_{it} \approx \theta_{Mi} \Delta \ln P_{M,t} + \theta_{Fi} \Delta \ln F_t$, where θ_{Mi} and θ_{Fi} are product–level output elasticities. Combining these ingredients with the linear pricing response in (E1) yields the estimating equation

$$\Delta p_{it} = \alpha(\theta_{Mi} \Delta \ln P_{M,t} + \theta_{Fi} \Delta \ln F_t) + (-\beta_s \phi_i) S_{it} + \varepsilon_{it}, \quad (\text{E2})$$

so the coefficients on the cost series recover $\alpha\theta_{Mi}$ and $\alpha\theta_{Fi}$, while the shortfall coefficient $(-\beta_s \phi_i)$ is larger exactly where import flows are a more important source of inventory.¹⁸ In what follows, we abstract from firm-specific heterogeneity in θ_{Mi} , θ_{Fi} , and ϕ_i , and interpret the estimated coefficients as average effects across firms.

3 Data and Measurement

We construct a large-scale micro dataset that links consumer prices to firms’ import activity in order to quantify how delivery shortfalls and cost shocks pass through to retail prices. We combine two complementary sources: a detailed product-level consumer panel from Numerator,¹⁹ which provides transaction-level prices and sales, and shipment-level Bills of Lading (BoL) data from Panjiva, which cover nearly the universe of U.S. imports and contain detailed economic and logistical information on each shipment. This linkage allows us to map shocks to imported inputs and shipping conditions into observed product-level price changes.

¹⁸The precise construction of S_{it} is deferred to the data section.

¹⁹We thank Leo Feler and the data team at Numerator for providing us with data access and invaluable support.

Guided by the model, our empirical design separates marginal-cost and availability (stock-in) channels. Accordingly, we construct three sets of supply-chain measures. First, *cost shifters*: HS-level unit import values and port-level freight/insurance ratios, aggregated to the firm using predetermined (2019) import-exposure weights, proxy shifts in imported-input and transport components of marginal cost. Second, *availability*: firm-level delivery shortfalls—measured as deviations from the firm’s pre-period shipping schedule—proxy stock-in probability and capture the non-price effects of delays on pricing incentives. Third, *congestion instruments* for firm-level delivery shortfalls: exposure to HS-level shortfalls (leave-one-out) and port dwell-time changes shift a firm’s predicted shortfall through supply-chain congestion and are used to isolate the supply-driven component of availability changes, conditional on detailed product market-time fixed effects. Below, we first introduce the consumer panel and the Bill of Lading data, and then describe how we construct each series in turn.

3.1 Numerator Data and Price Measurement

We use consumer panel data from Numerator for the period 2019-2023 to infer prices and sales of various products by firms.²⁰ Numerator is a marketing research company that collects data on consumer purchases in brick-and-mortar stores and online. Participants in the Numerator app either upload photos of their receipts or passively share purchase information by allowing the app to track online transactions and loyalty program activity. The closest counterpart to these data is the widely used Nielsen Homescan (HMS) data from the Kilts Center.²¹ For our purposes, Numerator’s advantages over HMS include its coverage of the recent pandemic period and its inclusion of e-commerce, whose share of total retail sales has substantially increased since the pandemic.²² Appendix Section B.1 provides more details about the data.

The Numerator dataset contains over one billion receipts from 3.9 million users, forming the basis of a *static panel* of 450,000 panelists (200,000 monthly) who report purchases for at least 12 consecutive months and are representative of the U.S. population along key Census demographics. For each static panelist, the data include demographic weights (*demo_weight*) for alignment with gender, age, ethnicity, household income, household size, presence of children, and census division, and projection factors (*national_factor* and *trend_factor*) to scale from sample households to the U.S. population. We use the combined weights and factors in all analyses of quantities and sales.

²⁰Analysis is based on the April 19, 2024, Numerator Data Delivery.

²¹The Numerator data have only recently begun to be used in academic research, having been described and benchmarked against similar datasets in prior studies (He and Su, 2023; Hacıoğlu Hoke, Feler and Chylak, 2024; Hristakeva, Liaukonyte and Feler, 2024) and more recently applied in Baslandze et al. (2025).

²²Approximately 90% of the transactions are from offline purchases, while the remaining transactions come from digital sites such as Amazon.com.

Each shopping receipt recorded in the Numerator data provides detailed information about the purchased basket, including item descriptions, quantities purchased, prices, and details about the stores (such as name and address) where the purchases occurred. Important for our analysis is the information on the purchased items. Where available, Numerator item IDs are linked with UPC/GTIN codes, brands, and producer names.²³ Items are classified into different categories with varying levels of aggregation. We exclude purchases that lack classification or are categorized as restaurants, non-items, or intermediate categories, focusing on item IDs with non-missing brand or manufacturer information.

Comparisons of Numerator data with official statistics suggest that the data provide significant coverage of product sales and consumer behavior in the U.S. First, the demographics of static panelists are similar to those reported in the 2019 Census. Appendix Figure A.1 compares the demographics of static panelists with the U.S. Census and shows that the panelist sample closely resembles the representative U.S. household in terms of age, region, education, employment, ethnicity, gender, number of children, marital status, and income (even without applying demographic weights).

Second, products covered by Numerator are representative of goods and services categories accounting for approximately 40% of consumer expenditures in the CEX.²⁴ Numerator data exclude consumer expenditures on shelter and utilities (26% of CEX), insurance and financial services (11.9%), a large portion of transportation (6.2%), and healthcare (8%). The major consumption categories with the highest expenditure shares in Numerator are grocery (43% of Numerator spending), health and beauty, electronics, household products, apparel, restaurants, and others (see Appendix Figure A.2).

Third, changes in total sales across different purchase categories (such as apparel, electronics, health and beauty, and food) closely track the sales changes observed in U.S. Census Monthly Retail Sales data (Numerator, 2023). Most importantly for our analysis, monthly price changes in Numerator closely align with monthly price changes reported by the BLS for various purchase categories, as shown below.

For our baseline analysis, we aggregate 54 million item IDs reported in Numerator and define a *product* as the unique combination of the firm, brand, and product category. For example, a product might be a Hasbro (firm) producing a Peppa Pig (brand) 3-D puzzle (product category). This approach yields 645,890 distinct products and 77,395 brands sold by 9,458 firms across 4,000 product categories.²⁵

For each item and month, we compute the average price as weighted total sales divided by weighted total quantity, using all transactions from static panelists. *Product-level price* $Price_{p,t}$ — with p being a product and t being a month— is then obtained as a quantity-weighted average of item-level prices.

We construct the *product-level price change* as a 12-month smoothed price change for product p in

²³Many products or services are not associated with UPC codes, such as restaurant or entertainment purchases.

²⁴Consumer Expenditure Surveys, 2022, Bureau of Labor Statistics.

²⁵Other examples of products include Hewlett Packard/Hewlett Packard/Calculators, Revlon Consumer Products/Revlon/Flat Irons, American Italian Pasta Company/Mueller's/Lasagna, and Loreal/Cerave/Bar Soaps.

month t , defined as:

$$\Delta_{12} \text{ Price}_{p,t} = \frac{\text{Price}_{p,t} - \text{Price}_{p,t-12}}{(\text{Price}_{p,t} + \text{Price}_{p,t-12})/2}.$$


This measure, inspired by [Davis, Haltiwanger and Schuh \(1998\)](#), captures relative price changes over a one-year horizon and reduces the influence of outliers, allowing for robust analysis of pricing dynamics.²⁶ To further reduce outliers and account for data noise, we trim the top and bottom 1% of 3, 6, and 12-month price changes (including missings) and drop all products with fewer than five observations over time and having fewer than five receipts collected in a month.

To validate that our micro-level price data capture aggregate inflation dynamics, we compare inflation rates constructed from Numerator to official CPI measures. Numerator provides a reliable measure of inflation for the consumer goods we study and captures the broader inflationary pressures that emerged during the pandemic. Figure 3 shows that 12-month inflation rates constructed from Numerator closely track the Urban CPI, with both series peaking in mid-2022 and exhibiting similar dynamics throughout the period. Although Numerator primarily covers consumer goods, food, and grocery items—and excludes major CPI components such as shelter and services—the strong comovement suggests that goods-sector inflation accounts for a substantial share of overall inflation dynamics during this episode. Appendix Figure A.3 provides a more direct comparison for food categories, showing a tight match between food CPI and price changes constructed from food and grocery items in Numerator. The CPI series appears somewhat smoother, likely reflecting differences in aggregation and seasonal adjustment, but the underlying trends are highly similar.

Finally, the availability of micro-level product price data allows us to document substantial heterogeneity in price changes over time—both across categories and, within narrowly defined categories, across firms—particularly during the pandemic (Appendix Figure A.4). This dispersion provides the key product-level variation that underpins our identification strategy. Even after controlling for detailed category and time fixed effects, a large share of price variation remains at the product–firm level. Specifically, regressions of product-level price growth on category and time fixed effects yield R^2 values below 10 percent. Thus, common sectoral or aggregate shocks explain only a small fraction of observed price movements. The overwhelming majority reflects within-category, cross-firm differences, providing the granular variation that underpins our identification in the analysis that follows.

²⁶We verify the robustness of our results by comparing this smoothed measure with an alternative definition in log changes, confirming that our findings are consistent across different specifications.

Figure 3. Numerator Price Change vs. CPI



Notes: The solid red line plots the 12-month aggregate price change derived from Numerator transaction data. Product-level price changes are first averaged within categories and then aggregated using category sales weights. The dashed blue line plots U.S. CPI inflation for all items (12-month percent change; FRED series CPI-AUCSL). The close co-movement indicates that the Numerator panel captures broad inflation dynamics over the sample period.

3.2 Panjiva Bill of Lading Data (BoL)

The analysis relies on bill of lading (BoL) data to infer import quantities and delivery shortfalls. A BoL is a legal document between the shipper and carrier required for the shipment of goods, detailing the commodity, shipping and receiving entities, vessel identifier, port of entry, payment terms, and other logistics details. These data are derived from images of contracts processed from U.S. ports and customs, providing a unique, high-frequency view of trade relationships at the transaction and firm level. BoL data are particularly valuable for capturing detailed maritime trade and logistics information in a timely manner, making them well-suited for analyzing international delivery delays.

We use bill-of-lading (BoL) data from S&P Panjiva, which provide comprehensive records of U.S. maritime imports since 2007, covering over a billion shipments (see [Flaaen et al. \(2023\)](#) for a detailed description, comparisons with Census data, advantages over alternative sources, and data limitations). BoL data are increasingly used as a publicly accessible and high-frequency source of U.S. import information. An important advantage is the availability of consignee and shipper names, which enables matching to a wide range of firm-level datasets.

At the same time, the data have several limitations. First, they cover only seaborne imports. In 2020, for example, maritime transport accounted for 43% of U.S. imports by value and 60% by weight (Bureau of Transportation Statistics.²⁷ Second, while BoL data provide reliable information on quantities—an important input in our shortfall measures—the reported import values, which we use to construct import shares, are often imputed and less reliable. To address this issue, we use public Census data from

²⁷Since these are maritime shipments, coverage from Mexico and Canada is very limited.

U.S. Trade Online to convert BoL import volumes (in TEU) into values. Specifically, we compute HS2-specific 2019 volume-to-value ratios and apply them to BoL shipment volumes at the firm–HS2–month level. Finally, the data contain some firm-name redactions; we mitigate this issue by focusing on stable importing firms to reduce noise, as further discussed below. Appendix Section B.2 provides additional details on the BoL data, the cleaning procedures, and validation against the U.S. Census series.

Our final dataset covers U.S. imports from 2019–2023, aggregated to a firm–HS2–month panel that includes shipment volumes and values, baseline (2019) benchmarks, and moving-cumulative measures of import volumes. These data allow us to measure deviations from typical shipment volumes to detect delivery delays, which we describe below. In addition, we combine these data with port congestion and shipping delay indicators derived from Automatic Identification System (AIS) data. These complementary sources, which capture vessel trajectories and port dwell times, help distinguish between delays arising from port congestion and those specific to the shipment’s origin.

3.3 Delivery Shortfalls, Instruments, and Cost Shocks

We describe the construction of (i) firm-level delivery shortfalls, which proxy for firm-level product availability; (ii) two instruments for these shortfalls—a shift–share Delivery Shortfall Exposure measure and a port-congestion exposure measure; and (iii) two import-based cost shifters: unit import costs and freight and insurance cost shocks.

Delivery Shortfalls

Inspired by our discussion in Section 2, the core measure of interest, delivery shortfall, is constructed to quantify deviations in k -month cumulative import volumes relative to a pre-pandemic baseline, capturing the impact of delivery delays at the firm and product level. Specifically, monthly delivery shortfalls are calculated as a smoothed percent deviation of k -month cumulative imports in a given month from their 2019 levels in the same calendar month. Hence, the year 2019 serves as a pre-disruption benchmark; and deviations from that benchmark “delivery schedule” proxy for inventory disruptions at the firm.

Starting from the delivery shortfall at the firm–HS-code level, we define:

$$\text{Delivery Shortfall}_{fht}^{(k)} = -\frac{\sum_{j=0}^k D_{fh,t-j} - \sum_{j=0}^k D_{fh,t^{2019}-j}}{\left(\sum_{j=0}^k D_{fh,t-j} + \sum_{j=0}^k D_{fh,t^{2019}-j} \right) / 2},$$

where t denotes time (year–month), k is the chosen horizon of cumulation, and $D_{fh,t}$ represents the observed import volume for firm f and product code (HS) h in month t . The term t^{2019} denotes the month in 2019 with the same calendar month as t . The leading minus sign ensures that the measure is positive when realized deliveries fall short of the 2019 schedule.

This firm-HS-code-level measure is then aggregated to the firm level by weighting the HS-code shortfalls by their shares in the firm’s total imports in the pre-disruption baseline year, 2019:

$$\text{Delivery Shortfall}_{ft}^{(k)} = \sum_h \omega_{fh}^{2019} \times \text{Delivery Shortfall}_{fht}^{(k)}, \quad (1)$$

where ω_{fh}^{2019} denotes the share of each HS-code h in the total imports of firm f in 2019. This aggregation captures the overall impact of delivery disruptions at the firm level, accounting for the relative importance of each product code in the firm’s baseline portfolio.

Our delivery shortfall measures are constructed using a stable sample of firms that had some import activity in 2019 and maintained an import presence in each subsequent year. This approach reduces data noise (including issues related to firm-name redactions in the data) and focuses the analysis on firms that consistently rely on imports.²⁸ For this sample, we identify each firm’s portfolio of HS codes and their import schedules in 2019 and track firm-HS-code (k -month cumulative) deliveries over time (months) relative to the benchmark in 2019 (including zeros).²⁹ Finally, we normalize the shortfall series to January 2020.

For our main analysis, we construct firm import portfolios and shortfalls at the two-digit HS (HS2) level. This choice primarily reflects measurement considerations. HS codes in BoL data can be noisy or inconsistently assigned at finer levels,³⁰ and shortfall construction at the monthly firm–HS level becomes sparse and volatile when using more disaggregated codes. Aggregating to HS2 delivers more stable and reliable monthly series, improves signal-to-noise in the shortfall measures, and strengthens identification in the shift–share design, as further discussed below. While HS2 is coarser and bundles heterogeneous goods, we interpret any within-bin composition changes as measurement error that would tend to attenuate our estimates.³¹


²⁸There is substantial entry and exit in the BoL data, as shown in Appendix Figure A.6, which plots total import volumes for the sample of firms with any deliveries in 2019—comparable to Figure A.5. While focusing on a stable sample of importers helps reduce data noise, it also drops those firms that experienced extreme negative events and exited the market, potentially leading to underestimating the negative effects of supply chain disruptions.

²⁹This implies that entry of new HS codes within a firm over time obtains zero weights in these calculations, so that we trace a deviation relative to the 2019 benchmark.

³⁰HS codes in Panjiva are not official customs classifications but are largely assigned by the data provider (Flaen et al., 2023). On average, HS codes are directly parsed from the bill of lading in about 20% of observations, while in the remaining 80% they are assigned algorithmically, with some variation in these shares over time.

³¹For completeness, we also reconstruct the series at the more disaggregated HS4 level and replicate our main specifications. The core results remain largely robust across aggregation levels, although, as expected, the HS4 series are noisier. See

Figure 4. Average Delivery Shortfall Over Time



Notes: The figure plots the weighted average firm-level delivery shortfall over time using the Panjiva sample. Shortfalls are measured relative to the firm’s 2019 baseline import schedule. Firms are weighted by total annual import volume.

Figure 4 shows the evolution of average firm-level delivery shortfalls using our baseline horizon of $k = 3$, along with alternative windows ($k = 2$ and $k = 6$). These shortfalls capture supply chain disruptions, with larger positive deviations indicating more severe delays in the flow of goods. A sharp increase in 2020, relative to the 2019 benchmark, reflects the unprecedented disruptions triggered by pandemic-related lockdowns and labor shortages. Shortfalls rebounded in 2021, though the series continues to exhibit fluctuations. As expected, longer moving averages smooth out short-term volatility and result in less variable trends.

The shortfall measure—exploiting the granularity of BoL data—captures deviations driven by shipping disruptions, production delays, and other logistical frictions, providing a direct reflection of firm-level supply chain shocks. However, in its raw form, changes in delivery schedules may also reflect aggregate or firm-specific demand factors. To isolate the supply-driven component of the shortfall, our empirical specifications rely on detailed fixed effects and an instrumental variables (IV) approach to filter out aggregate and firm-level demand forces. We now turn to the construction of these instruments.

Instruments for Delivery Shortfalls

Instrument 1: Delivery Shortfall Exposure. Our first instrument for the firm-level delivery shortfall is a standard shift-share measure capturing a firm’s exposure to aggregate delivery disruptions across

Appendix Table A.7.

its import portfolio:

$$\text{Delivery Shortfall Exposure}_{ft}^{(k)} \equiv \sum_{h \in S^{HS2}} \omega_{fh}^{2019} \times \text{Delivery Shortfall}_{ht,-f}^{(k)}, \quad (2)$$

where $h \in S^{HS2}$ indexes HS2 product codes, ω_{fh}^{2019} is the predetermined share of firm f 's imports in HS code h in 2019, as before, and $\text{Delivery Shortfall}_{ht,-f}^{(k)}$ denotes the HS-code-level aggregate shortfall in month t relative to the same month in 2019, constructed excluding the focal firm f itself (leave-out shifts).


Figure A.7 illustrates the sources of variation underlying our shift-share exposure measure. Panel A focuses on the “shifts.” It shows the distribution of aggregate HS-code-level shortfalls, defined as deviations of imports from their 2019 baseline (constructed analogously to the shifts, though not in leave-out form). The figures reveal substantial heterogeneity in how aggregate deliveries evolved across HS codes over time, providing meaningful variation in the shift component.

Turning to the “shares”—measured by firms’ baseline import shares across HS2 categories in 2019—we also observe considerable variation in sourcing patterns. Firms differ both in the set of HS codes they import and in the intensity with which they rely on each code. On average, firms import from 5.8 HS2 codes, and the average firm-level HHI of 0.66 indicates that firms typically combine multiple inputs rather than relying on a single code.

Panel B further documents dispersion in import shares across firms within the same (Numerator-based) product category. For each HS2 code, we compute the standard deviation of import shares across firms within a category and then average this dispersion across categories. The figure shows substantial variation for most HS2 codes, indicating that firms selling similar products often rely on different import structures. Consistent with this, product-market fixed effects explain only a small fraction of the variation in import shares. In a regression of firm-HS2 import shares on category fixed effects, the R^2 is about 5 percent, implying that most of the variation occurs across firms within the same product category. Together, these results confirm that both the shift and share components of the shift-share design exhibit substantial and economically meaningful heterogeneity.

Instrument 2: Port Congestion Exposure. Our second instrument is a port congestion exposure using dwell-time variation across ports. The idea is that the longer the dwell times due to port congestion, the more likely a delivery shortfall is to occur. This instrument leverages exogenous variation in port-level congestion, derived from detailed port dwell time data, combined with firm-specific historical port usage patterns. The port dwell time data itself has been constructed from high-frequency Automatic Identification System (AIS) vessel traffic information which provides high-frequency location

Figure 5. Port Congestion Exposure: Change in Dwell-Time Exposure Over Time



Notes: The figure plots the weighted average change in firm-level port dwell-time exposure for Panjiva firms. Exposure is constructed as a shift-share measure using fixed 2019 port shares interacted with residualized port-level dwell-time changes (see Section 3.2). This measure captures congestion-induced delivery delays and is used to isolate supply-driven variation in firm shortfalls in the IV strategy. Firms are weighted by total annual import volume.

information of individual vessels at U.S. ports (Fuchs and Wong, 2022). Dwell time is then defined as the duration a vessel spends moored at a pier with zero speed. Crucially, to isolate port-specific conditions from variations due to the mix of ships calling, these raw dwell times (in logs) are residualized by controlling for ship characteristics, particularly size (e.g., gross tonnage). This yields a measure of port-level delay (in logs) (Residualized Dwell Time_{pt}) for port p at time (year-month) t that is net of ship-specific factors.

To construct firm-level port-congestion exposure measure, we combine this port-level “shift” measure with firm-specific “shares” derived from BoL data. The BoL data provides granular information on import shipments at the firm level, including the port of entry. From this, we calculate each firm f ’s historical reliance on different ports, measured as the share of its total imports that arrived through a specific port p during a pre-determined base period (ω_{fp}^{2019}). Using 2019 as a base period ensures these shares reflect established logistical relationships. The firm-level shift-share instrument, representing firm f ’s exposure to port congestion, is then constructed by weighting the port-specific residualized dwell time shifts by the firm’s historical port shares, summing across all ports:

$$\Delta \text{Dwell Exposure}_{ft} \equiv \sum_p \omega_{fp}^{2019} \times \Delta \text{Residualized Dwell Time}_{pt}. \quad (3)$$

This instrument is designed to capture the variation in a firm’s exposure to import delays driven by congestion at the specific ports it historically uses, isolating it from firm-specific demand or supply shocks. The validity of the instrument relies on the standard shift-share assumptions: relevance (the

instrument predicts actual firm-level delays/shortfalls) and exclusion (the instrument, conditional on controls, only affects the firm’s outcomes through its impact on import delays/congestion). That is, port-level congestion patterns from ports a firm historically used should not directly affect the firm’s outcomes today, other than via the channel of current import disruptions.

Figure 5 plots the evolution of the average change in firm-level dwell time exposure. The series shows a sharp increase in port congestion during the heightened inflationary period, reflecting severe bottlenecks in global shipping and longer delays in moving goods through US ports.

In summary, the two shift-share instruments isolate the component of firm-level delivery shortfalls that is driven by common supply-chain disruptions rather than firms’ own ordering or demand conditions. The first instrument, *Delivery Shortfall Exposure*, loads a firm’s pre-pandemic HS2 import mix on HS2-level shortfalls experienced by other importers, capturing upstream and product-specific disruptions in the goods the firm sources. The second instrument, *Port Congestion Exposure*, loads a firm’s pre-pandemic port mix on residualized port dwell-time spikes, capturing congestion-induced delays at the ports the firm historically uses. Because the exposure shares are fixed at the baseline and we include market×time fixed effects, identification comes from comparing firms selling in the same product market at the same time but with different pre-period reliance on disrupted HS2 categories and congested ports.

Import-Based Cost Shifters: Unit and Freight Costs

We construct two import-based cost shifters that isolate changes in input costs arising from unit import prices and freight costs. Using U.S. Census merchandise-trade data ([U.S. Census Bureau, 2026](#)), we first derive HS-level changes in unit import values and port-level freight and insurance ratios (“shifts”). We then aggregate these common HS/port series to the firm using predetermined (2019) Panjiva import-portfolio shares.³² Specifically, we construct: (i) a *unit import cost* index at the HS2 level (indexed by h), $UV_{ht} \equiv FOB_{ht}/Q_{ht}$ (customs value per physical quantity), and (ii) a *freight/insurance* index at the port level (indexed by p), $FR_{pt} \equiv CIF_{pt}/FOB_{pt}$. We take 12-month log changes and map these series to firm f using fixed pre-period (2019) exposure weights derived from Panjiva:

$$\Delta \text{UnitC}_{ft} = \sum_h \omega_{fh}^{2019} \Delta_{12} \ln UV_{ht}, \quad \Delta \text{FreightC}_{ft} = \sum_p \omega_{fp}^{2019} \Delta_{12} \ln FR_{pt}, \quad (4)$$

³²Because unit values and freight ratios are aggregate *average* values and do not use firm quantities, we interpret them as shifters to the imported-input and shipping components of costs, not as firm-level marginal cost measures. Any mismatch between these shifters and the marginal unit relevant for pricing appears as proxy error and is most naturally expected to attenuate pass-through estimates toward zero.

where ω_{fh}^{2019} is the share of f 's 2019 imports in HS2 code h , and ω_{fp}^{2019} is the share from port p .³³

Figure 6 plots the evolution of the cost shifters: exposure to changes in unit import values and freight/insurance ratios over time. Both components of import costs rose sharply during the inflationary period, reflecting surging global input prices and severe pressures in shipping and logistics. At the same time, comparing Figures 5 and 6 makes clear that the supply-chain objects we use exhibit distinct time profiles during the pandemic. Port congestion (dwell-time exposure) increases sharply in 2021–2022 and then recedes; freight/insurance ratios spike and unwind on a different schedule; and HS-level unit values rise more gradually.

These differences in timing motivate allowing for short-run dynamics in pass-through rather than imposing purely contemporaneous effects. In the baseline regressions, we therefore include both contemporaneous and lagged cost shifters. Delivery shortfalls proxy multi-month cumulative disruptions to availability and are expected to load primarily on the current term, whereas imported-input and freight cost shifters may transmit with a short delay due to inventories, contracts, and staggered price adjustment.


In summary, our two import-based cost shifters are designed to capture common movements in imported-input prices and shipping charges that plausibly shift firms' marginal costs but are not mechanically tied to firm-specific pricing decisions. The *Unit Import Cost* shifter loads a firm's pre-determined HS2 import mix on HS2-level unit import value changes, capturing upstream price pressure in the inputs it typically sources. The *Freight Cost* shifter loads a firm's pre-determined port mix on port-level freight/insurance ratio changes, capturing time-varying shipping charges along the firm's usual logistics routes. Because these cost series vary at the HS- and port level and are mapped to firms using pre-period shares, the identifying variation comes from comparing firms in the same product market and time period that differ in their pre-pandemic exposure to input-cost changes and shipping-cost spikes.

3.4 Merging Consumer Prices with Imports Data: Numerator-Panjiva Match

To identify importing firms in our consumer product data and relate supply chain disruptions to prices, we link Numerator consumer purchase data to BoL records from Panjiva. The match is constructed using firm names across datasets. Starting from the manufacturer, parent brand, and brand fields in Numerator, we hierarchically match to consignee names (and, when necessary, shipper names) in Panjiva. This procedure relies on extensive name cleaning and standardization to address spelling

³³For exposition, one can also summarize the imported-input contribution as $\Delta MC_{f,t} \approx \phi^c(\Delta \text{Unit}C_{f,t} + \Delta \text{Freight}C_{f,t})$; in practice we use the two components separately.

Figure 6. Import Cost Shifters: Unit Values and Freight-Cost Ratios



Notes: Panel (A) plots the weighted average firm-level exposure to HS2 unit import value changes, $\Delta UnitC_{ft}$. Panel (B) plots the weighted average exposure to port-level freight/insurance ratio changes, $\Delta FreightC_{ft}$ (import charges relative to customs value). Firms are weighted by total annual import volume.

differences, formatting inconsistencies, and embedded location details in Panjiva’s name fields.³⁴

The resulting firm-level linkage allows us to split Numerator firms into importing and non-importing groups. Of the 3.5 million product-month observations in our cleaned baseline Numerator sample (2019–2023), 2.6 million correspond to firms that match at least once to Panjiva BoL records. Among these, 1.5 million observations also merge to our cleaned baseline sample of stable importers for which we construct the shortfall measure (described above). We refer to this intersection of the baseline Numerator and stable-importer Panjiva samples as the baseline Numerator–Panjiva sample.

Table 1 reports summary statistics for the full Numerator sample and the matched baseline sample. The latter accounts for 44% of product-month observations and 51% of total sales, even though importing firms represent less than 15% of distinct firms in the data. Thus, a relatively small set of importers generates roughly half of observed consumer sales. Consistent with this concentration, importing firms are substantially larger on average, both in total sales and in the number of transactions.

Appendix Table A.3 compares importing and non-importing firms more directly. The two groups operate in largely overlapping product markets, spanning almost all departments covered in Numerator. Firms classified as non-importing are those who do not actively import in Panjiva. Because Panjiva primarily captures seaborne trade, some of these firms may still engage in international trade via air or other channels. Moreover, even firms that do not import directly may be indirectly exposed to foreign supply shocks through domestic input–output linkages. For these reasons, our importer classification likely understates total foreign exposure, implying that our estimated supply-chain effects should be interpreted as conservative lower bounds.

³⁴Appendix B.3 provides more details on the cleaning steps and matching procedure.

Table 1. Summary Statistics at the Product-Month Level

	Numerator Sample (All Firms)	Numerator–Panjiva Sample (Importing Firms)
Number of Observations	3,501,353	1,542,111
Date Range	2019–2023	2019–2023
Distinct Number of Products	130,459	50,375
Distinct Number of Firms	29,444	4,174
Distinct Number of Product-Categories	3,987	3,633
Distinct Number of Product-Departments	217	213
Average Price Change (12-month)	0.039	0.041
Average Sales	1,380,073	1,713,221
Average Transactions (All Users)	700	915
Average Transactions (Static Users)	346	439
Panjiva Firm Dummy	0.440	1.000

Notes: Summary statistics are reported at the product-month level for 2019–2023. Column (1) corresponds to the full cleaned Numerator sample (see Section 3.1). Column (2) reports the baseline Numerator–Panjiva matched sample of stable importing firms (see Section 3.2). The matched sample accounts for approximately 44% of product-month observations and 51% of total sales.

Lastly, the merged microdata on prices and imports reveals how price changes during the recent inflationary episode vary across firms. Interestingly, Figure A.8 shows that larger firms and those with higher market shares experienced larger 12-month price increases, conditional on product-department and year fixed effects.³⁵ At the same time, price growth is negatively correlated with measures of sourcing diversification: firms that import from a larger number of countries (either overall or within specific HS products) tend to exhibit smaller price increases, conditional on size. Together, these raw correlations suggest that larger firms raised prices more, and those with more diversified sourcing networks were better insulated from supply disruptions.

4 Analysis

This section examines how supply-chain disruptions—delivery shortfalls, unit import cost shocks, and freight cost shocks—affect firm pricing. We first estimate the pass-through of a firm’s own disruptions, characterize the timing of price adjustments, and document heterogeneity across product categories. We then introduce strategic interactions and show that competitors’ disruptions also raise a

³⁵These patterns are consistent with Franzoni, Giannetti and Tubaldi (2023), who document that larger firms expanded market shares and markups relative to smaller competitors during the pandemic inflationary episode in the United States.

firm's prices—over and above the effect of its own shocks, and that these spillovers extend even to non-importing firms.

4.1 Pass-Through of Own Supply Chain Disruptions to Prices

Empirical Specification and Identification

Guided by the simple model's local decomposition $dp = \alpha dmc + \beta_s d \ln s$, we estimate a reduced form in which (i) availability enters via firm-level delivery shortfalls, our proxy for $\Delta \ln s$, and (ii) marginal cost change is decomposed into unit-import and freight components. Concretely, we estimate

$$\Delta \text{Price}_{pft} = \phi_s \text{Shortfall}_{ft}^{(k)} + \phi_M \Delta \text{UnitC}_{ft} + \phi_F \Delta \text{FreightC}_{ft} + \theta_f + \theta_{j(p)q(t)} + \epsilon_{pft}, \quad (5)$$

where $\Delta \text{Price}_{pft}$ is the 12-month change in the price of product p at firm f in month t ; $\text{Shortfall}_{ft}^{(k)}$ is the k -month cumulative delivery shortfall (baseline $k = 3$, Section 3.1); and ΔUnitC_{ft} , $\Delta \text{FreightC}_{ft}$ are firm shift-share exposures to unit-import and freight cost changes (Section 3.3). Firm fixed effects (θ_f) absorb time-invariant heterogeneity in pricing, and category $\times$ quarter fixed effects ($\theta_{j(p)q(t)}$) absorb common market-specific demand/supply movements.

Equation (5) is the direct empirical counterpart to the model's local pricing decomposition in (E2): we allow consumer price changes to load separately on (i) imported-input and freight cost *shifters* and (ii) *availability*, proxied by delivery shortfalls. On the availability side, delivery shortfalls are a monotone proxy for stock-in probability, $\Delta \ln s_{f,t} \approx -\phi \text{Shortfall}_{f,t}^{(k)}$, with $\phi > 0$, so the estimated shortfall coefficient recovers the availability channel scaled by measurement, $\phi_s = -\beta_s \phi$, and is therefore expected to be positive. On the cost side, the model implies $\Delta mc_{ft} \approx \theta_M \Delta \ln P_{Mt} + \theta_F \Delta \ln F_t$, so ϕ_M and ϕ_F identify $\alpha \theta_M$ and $\alpha \theta_F$, respectively. For comparability with the cost pass-through literature, we scale right-hand-side cost exposures by industry import-intensity weights and scale shortfalls by intermediates import shares.³⁶

Idiosyncratic demand shocks may bias OLS estimates of the shortfall coefficient ϕ_s . For example, a positive firm-specific demand shock can prompt a firm to accelerate orders, mechanically shrinking measured shortfalls and attenuating ϕ_s toward zero. Measurement error in the shortfall proxy (our local linear map from $\Delta \ln s$ to $\text{Shortfall}_{f,t}^{(k)}$) can generate the same attenuation. Hence, OLS estimates of ϕ_s are likely biased toward zero.³⁷ We therefore implement a shift-share IV strategy that

³⁶Using BEA Input–Output Accounts, we use an import share of 0.094 for industries mapping to Numerator consumer-product sectors to scale the unit-cost and freight-cost series, and an intermediates import share of 0.125 to scale shortfall-related series.

³⁷Two forces make us expect the OLS estimate of the shortfall coefficient ϕ_s to be biased downward. First, shortfalls

instruments $\text{Shortfall}_{f,t}^{(k)}$ with $\text{Delivery Shortfall Exposure}_{f,t}^{(k)}$ and $\Delta \text{Dwell Exposure}_{f,t}$, as described in (2) and (3). Conditional on firm fixed effects and product-category $\times$ quarter fixed effects, identification comes from comparing firms selling in the same product market at the same time but with different pre-pandemic sourcing patterns—whether due to reliance on different HS codes or importing ports. These instruments, therefore, isolate variation in delivery shortfalls driven by pre-determined exposure to common supply shocks rather than firm-specific demand conditions. In practice, the instruments are strong, with first-stage F-statistics well above conventional thresholds.

Under an exogenous-shares (parallel-trends) interpretation, the shifts may reflect equilibrium movements and need not be exogenous on their own. Identification instead requires that, conditional on controls, firms with higher 2019 exposure to a given HS code or port would have experienced similar within-market price trends absent the pandemic disruptions.³⁸ Under this assumption, the exposure measures provide firm–time variation in predicted shortfalls that is orthogonal to idiosyncratic demand shocks, and the second stage recovers the supply-driven component of availability pass-through.

Several considerations are important when interpreting the pass-through estimates. First, our empirical cost objects—HS-level unit values and port-level freight/insurance ratios—are *aggregate average* price indices and therefore proxy (imperfectly) for the relevant marginal unit for firm pricing; we interpret them as *imported-input and freight cost shifters*, not literal firm-level marginal costs. Because the specification is in log changes, a broad class of convex procurement cost structures implies that log changes in average and marginal costs move proportionally up to a time-invariant wedge;³⁹ time-varying

are partly *endogenous* to firm-specific demand and ordering behavior. For example, a positive (unobserved) firm–product demand shock can raise prices while simultaneously inducing the firm to accelerate or expedite orders, mechanically reducing measured shortfalls relative to the 2019 shipping schedule. This generates negative comovement between $\text{Shortfall}_{f,t}$ and the pricing residual, attenuating the OLS estimate of ϕ_s toward zero (and, in extreme cases, potentially changing its sign). Second, delivery shortfalls are a noisy proxy for availability (e.g., HS aggregation and within-bin composition changes), which further attenuates the OLS coefficient through classical measurement error. Together, these considerations imply $\hat{\phi}_s^{IV} > \hat{\phi}_s^{OLS}$ if supply-driven disruptions indeed raise prices.

³⁸More formally, our IV strategy is a shift–share design in the sense of [Borusyak, Hull and Jaravel \(2025\)](#): the excluded instruments combine common disruption shocks at the HS and port level (the “shifts”) with predetermined dispersion in firms’ 2019 sourcing shares within narrowly defined product markets (the “shares”). Identification relies on heterogeneous exposure. Appendix Figure A.7 documents the necessary ingredients—heterogeneity in sourcing shares across firms and shifts of HS2 across time: Panel A shows substantial cross-HS and over-time dispersion in the aggregate shortfall shocks, and Panel B shows that firms selling within the same product category nonetheless have meaningful dispersion in their 2019 HS-sourcing shares, ensuring that the shift–share interaction generates nontrivial within-market variation (see also discussion in Section 3.3). Because shares are based on pre-period import portfolios (and are mechanically zero for non-importers), we include incomplete-share controls such as an importer indicator and, in robustness, pre-period import intensity. The parallel-trends/exogenous-shares assumption may break down if pre-period import portfolios correlate with unobserved firm characteristics that differentially interact with the pandemic (pricing power, domestic capacity constraints, demand reallocation, or other downstream cost changes) even within market $\times$ time cells.

³⁹In our empirical mapping the relevant object is $\Delta \log MC$. As in [Amiti, Itskhoki and Konings \(2019\)](#), log changes in average (unit) cost proxy log changes in marginal cost for a broad class of convex variable-cost schedules: if $TVC_t(Q) = C_t Q^{1+\nu}$, then $AC_t = C_t Q^\nu$ and $MC_t = (1 + \nu)C_t Q^\nu$, so $\Delta \log MC_t = \Delta \log AC_t$ for stable ν . Nevertheless, time-varying curvature or within-code mix changes show up as measurement error and attenuate pass-through estimates.

curvature and within-code mix/quality changes instead appear as measurement error that naturally attenuates estimated cost pass-through toward zero.

Second, although the simple model links movements in the P – MC wedge to markups, in the data firms combine imported goods with downstream inputs and services. As a result, the gap between consumer prices and our imported-input cost proxy reflects both markups and unobserved changes in other marginal-cost components, and we therefore avoid interpreting levels of P minus the cost proxy as markups. More generally, because our cost shifters capture only the import-related component of costs, they do not reflect other non-trade-related cost changes, such as domestic materials shortages, transportation bottlenecks, or capacity constraints. Accordingly, our estimates should be interpreted as the pass-through of *import-related* cost and availability shocks; broader domestic constraints are absorbed by the fixed effects when common within product-market–time cells, and otherwise remain in the residual.

Results

We begin by estimating baseline pass-through elasticities using both OLS and instrumental variables (IV) to identify the causal effect of supply chain disruptions—delivery shortfalls, import costs, and freight cost shocks—on prices. These estimates provide a benchmark for the magnitude and timing of firms’ price adjustments to these disruptions.

Table 2 reports the results corresponding to equation (5). The dependent variable is the 12-month price change at the product-month level, estimated on the Numerator–Panjiva matched sample for 2020–2023. Our baseline shortfall measure captures deviations from cumulative import deliveries over a three-month horizon ($k = 3$). Column (1) presents OLS estimates, while Column (2) replaces direct shortfalls with a shift-share shortfall exposure measure. Columns (3)–(5) instrument firm-level shortfalls using shortfall exposure and exposure to changes in dwell times. All specifications include firm fixed effects and category-time fixed effects to absorb constant firm-specific pricing factors and demand shifts at the product-category level. Table A.4 shows that the instruments are strong: the weak-IV F-statistics range from 181 to 395, well above conventional thresholds.

Across specifications, all pass-through coefficients on delivery shortfalls and cost shocks are positive, but OLS estimates on shortfalls are biased downward, as expected—likely reflecting unobserved demand shocks as discussed above. Using shift-share exposure in Column (2) raises the estimated effect, while IV estimation further increases the coefficients, yielding stable elasticities above 0.21. The pass-through elasticities on cost shocks are sizable but generally lower than the estimates in the literature.

Columns (4) and (5) extend the specification to include one-month lags of shortfalls and cost vari-

Table 2. Price Effects of Own Supply Chain Disruptions: Baseline Pass-Through Estimates


	(1)	(2)	(3)	(4)	(5)
	OLS	OLS (Shift-Share)	IV	IV	IV
Shortfall	0.006 (0.0047)	0.099*** (0.0373)	0.189** (0.0741)	0.224*** (0.0809)	0.249*** (0.0870)
$\Delta UnitC$	0.150** (0.0595)	0.153** (0.0594)	0.128** (0.0610)	0.023 (0.0657)	0.017 (0.0659)
$\Delta FreightC$	0.102** (0.0434)	0.109** (0.0435)	0.147*** (0.0479)	0.020 (0.0572)	0.027 (0.0580)
Lag Shortfall				0.036 (0.0767)	0.024 (0.0758)
Lag $\Delta UnitC$				0.189*** (0.0663)	0.193*** (0.0661)
Lag $\Delta FreightC$				0.234*** (0.0549)	0.233*** (0.0549)
Firm FE	✓	✓	✓	✓	✓
Category×Quarter FE	✓	✓	✓	✓	✓
Import dummy					✓
Observations	992,950	992,950	991,619	962,570	962,570
Weak IV F-stat			395.154	142.901	135.150

Notes: The table reports regressions of 12-month price changes on measures of own supply chain disruptions, using product-month Numerator–Panjiva matched data for 2020–2023. Column (1) estimates OLS using *Shortfall*. Column (2) replaces *Shortfall* with the shift–share *Shortfall Exposure*. Columns (3)–(5) report IV estimates that instrument *Shortfall* using *Shortfall Exposure* and *Dwell Time Exposure*. Columns (4)–(5) add one-period lags of the corresponding 12-month change/exposure series (e.g., $\Delta_{12} \ln UV_{h,t-1}$ mapped to the firm using 2019 shares). Column (5) additionally includes an import dummy equal to one if the firm has positive imports in that month. The baseline shortfall measure uses $k = 3$ (deviations from cumulative import deliveries over a three-month horizon). All specifications include firm and product-category×quarter fixed effects. Standard errors are clustered at the product-category×quarter level. ***, **, * denote significance at the 1%, 5%, and 10% levels.

ables. For shortfalls, only contemporaneous shocks matter; lagged effects are insignificant, consistent with the shortfall measure already reflecting cumulative three-month deviations. By contrast, for import and freight costs, lagged shocks are more important, suggesting that rising marginal costs pass through with a short delay. Column (5) also adds an import dummy (equal to one if the firm imports in that month), but this inclusion does not materially change the results.

We further investigate the timing of firms’ price responses to delivery shortfalls in Figure 7. The figure plots coefficients from separate IV regressions of 12-month product-level price changes on delivery

Figure 7. Pass-Through Estimates for Different Horizons of Delivery Shortfalls



Notes: The figure reports IV coefficients from separate regressions of 12-month product-level price changes on firm-level delivery shortfalls constructed using different moving-average horizons (in months). The specification follows Column (5) of Table 2 and includes current and lagged imported-input and freight cost shifters, an import dummy, and firm and product-category $\times$ quarter fixed effects. Points denote coefficient estimates and dashed lines represent 90% confidence intervals. Standard errors are clustered at the product-category $\times$ quarter level.

shortfalls defined at different horizons of moving average deviations (k months), following the specification in Column 5 of Table 2. The results show that the horizon of the delivery shortfall matters for pass-through. When disruptions are short-lived, firms tend to absorb more of the shocks. As shortfalls become more prolonged, however, pass-through to prices increases, suggesting that firms adjust once they recognize the shocks are persistent and cannot be absorbed internally.

Overall, the findings show that both disruptions to delivery schedules and increases in marginal costs due to higher import and freight expenses were significantly passed through to consumer prices during the pandemic.⁴⁰ Appendix Table A.5 examines heterogeneity by product type. While durable goods experienced markedly higher price growth during the pandemic, evidence of a higher elasticity of delivery-delay pass-through for durables is weak.⁴¹ This suggests that much of the observed price growth in durable goods reflected heightened demand rather than stronger supply-side pass-through.

In addition, we explore heterogeneity by firms' multinational status. Using Panjiva data, we construct a proxy for multinational firms based on transactions in which the consignee and shipper share the same name (e.g., intra-firm imports). Although these firms account for a substantial share of import value and product sales, we do not find systematic evidence of differential pass-through across

⁴⁰Our estimated elasticities reflect percent changes in retail prices with respect to border-based cost shifters. Because retail prices embed distribution costs and margins (e.g. Flaaen et al. (2025)), translating percent pass-through into dollar contributions requires additional assumptions on cost shares and distribution wedges. We therefore interpret our coefficients as percent pass-through rather than dollar incidence .

⁴¹Following Argente et al. (2020), we measure durability as the inverse of the frequency of trips the average household makes to purchase products in a department over a year. Table A.2 lists the most and least durable departments.

specifications.⁴²

4.2 Strategic Interactions

Extraordinary supply chain disruptions during and after the pandemic, together with an extreme and broad surge in inflation, led many commentators to wonder whether firms passed through aggregate shocks above and beyond the disruptions they directly experienced.⁴³ Yet firms’ pricing decisions reflect not only the pass-through of their shocks but also the market forces that govern consumer reallocation and discipline market power. Indeed, a sizable literature documents strategic complementarities in pricing across different settings (Amiti, Itskhoki and Konings, 2019; Albagli et al., 2025), naturally raising the question of whether consumer price increases during this period reflected such strategic interactions among firms. In this section, we extend our baseline analysis of own supply chain disruptions to explicitly incorporate competitive spillovers through rivals’ disruptions.

Theoretical motivation. As before, we start by motivating our empirical specifications with theory. Building on the simple model in Section 2, we now allow for strategic interactions across firms. Intuitively, each firm sets prices taking into account not only its own costs and availability but also how rivals’ conditions shift the overall competitive environment. Formally, a firm’s optimal price can be written as a fixed point in which markups depend on its own fundamentals and on the prevailing sectoral price and availability conditions.⁴⁴

The key step is to translate this best-response logic into an expression that depends only on fundamentals. Rivals’ prices, which appear in the firm’s first-order condition, can themselves be expressed in terms of their marginal costs and availability shocks. Aggregating those responses with share weights yields sector-level indices for competitors’ marginal costs and availability.⁴⁵ Substituting these rival

⁴²Multinational firms with intra-firm imports account for up to half of total import value in Panjiva and roughly 40% of products among importing firms in Numerator. The interaction between our shock measures and the multinational indicator is statistically insignificant in all specifications (results available upon request).

⁴³Baqae and Farhi (2022); Glover, Mustre-del Río and von Ende-Becker (2023). The idea of “greedflation” also gained traction, suggesting that firms exploited extraordinary shocks to raise prices to opportunistically increase profits – see Robert Reich’s testimony to Congress (April 5, 2022); Weber and Wasner (2023).

⁴⁴Formally, firm i ’s price satisfies $\tilde{p}_{it} = mc_{it} + \mathcal{M}_i(\tilde{p}_{it}, \tilde{\tau}_{it}, \mathbf{p}_{-it}, \boldsymbol{\tau}_{-it}; \boldsymbol{\xi}_i)$, with Γ_{it} and Λ_{it} denoting the own-price and own-availability derivatives, and Γ_{-it} and Λ_{-it} the sensitivities to rivals. Linearizing and substituting the firm’s own price reaction delivers the reduced-form above. Details in Appendix A.2.

⁴⁵We show in Appendix A that rivals’ aggregate responses can be written as $\Delta p_{-it} \approx \tilde{\alpha}_{-it} \Delta mc_{-it} + \tilde{\beta}_{-it} \Delta \tau_{-it}$, where $\tilde{\alpha}_{-it}$ and $\tilde{\beta}_{-it}$ are weighted averages of rivals’ pass-through and availability sensitivities.

responses back into firm i 's condition yields our empirical specification:

$$\Delta p_{it} = \underbrace{\frac{1}{1+\Gamma_{it}}}_{\alpha} \Delta mc_{it} + \underbrace{\frac{\Lambda_{it}}{1+\Gamma_{it}}}_{\beta} \Delta \tau_{it} + \underbrace{\frac{\Gamma_{-it}}{1+\Gamma_{it}} \bar{\alpha}_{-it}}_{\gamma^{mc}} \Delta mc_{-it} + \underbrace{\left(\frac{\Lambda_{-it}}{1+\Gamma_{it}} + \frac{\Gamma_{-it}}{1+\Gamma_{it}} \bar{\beta}_{-it} \right)}_{\delta^{\tau}} \Delta \tau_{-it} + \varepsilon_{it}. \quad (6)$$

The first two coefficients, α and β , correspond to the direct effects from Section 2: pass-through of own costs and the impact of own availability. The new terms capture *strategic interactions*. When competitors' costs rise, their prices adjust upward on average; because firm i 's markup condition depends on the sectoral price level, this generates a complementary response, captured by γ^{mc} . Rivals' availability, in turn, influences firm i both directly (more sectoral availability relaxes scarcity, pushing i 's price down) and indirectly through rivals' own price adjustments. The composite coefficient δ^{τ} therefore summarizes both a direct markup effect and an indirect price-mediated channel. The resulting specification, hence, nests the simple model's scarcity mechanism and quantifies how both own and rival supply-chain fundamentals transmit into prices through the same generalized-elasticity channel.

Empirical specification and identification. Building on these theoretical foundations, we extend the own-effects regression by adding competitor indices of delivery shortfalls and cost movements via unit import and freight cost shocks. Specifically, we estimate:

$$\begin{aligned} \Delta \text{Price}_{pft} &= \phi_s \text{Shortfall}_{ft}^{(k)} + \phi_M \Delta \text{UnitC}_{ft} + \phi_F \Delta \text{FreightC}_{ft} \\ &+ \psi_s \text{Shortfall}_{-fj(p)t}^{(k)} + \psi_M \Delta \text{UnitC}_{-fj(p)t} + \psi_F \Delta \text{FreightC}_{-fj(p)t} \\ &+ \theta_f + \theta_{j(p)q(t)} + \epsilon_{pft}. \end{aligned} \quad (7)$$

As before, $\Delta \text{Price}_{pft}$ is the 12-month price change for product p at firm f and month t , and the own variables are as in the baseline. Rival indices X_{-fjt} are leave-one-out, revenue-share-weighted averages of the rivals' variables— $X \in \{\text{Shortfall}, \Delta \text{UnitC}, \Delta \text{FreightC}\}$ —within market $j = j(p)$. Specifically, with firm revenue shares S_{gjt} for each firm g in market $j(p)$, define $\omega_{fgt} = S_{gjt}/(1 - S_{fjt})$ and set $X_{-fjt} = \sum_{g \neq f} \omega_{fgt} X_{gt}$. Firm fixed effects θ_f absorb time-invariant firm heterogeneity in pricing, and category-quarter effects $\theta_{j(p)q(t)}$ absorb common demand or supply shocks. As before, cost exposures and shortfall-related variables are scaled so that coefficients can be read as pass-through elasticities.

The competitor coefficients (ψ_s, ψ_M, ψ_F) are identified by comparing firms selling in the same market $j(p)$ at the same time, but facing different exposure to disrupted rivals as summarized by the leave-one-out, revenue-share-weighted indices X_{-fjt} . Category-quarter fixed effects absorb shocks common to all sellers in a product market over time, and firm fixed effects absorb time-invariant differences in pricing and product positioning. A natural concern is Manski's "correlated effects" (Manski, 1993): un-

observed shocks that differentially affect subsets of firms within a market (e.g., demand shifts across segments) could comove with both own prices and rivals' disruptions, potentially inflating ψ even absent strategic spillovers. Our specification mitigates mechanical reflection through the leave-one-out construction, and—on the availability side—by instrumenting own shortfalls with shift-share exposure to common HS/port disruptions as discussed above. We therefore interpret (ψ_s, ψ_M, ψ_F) as reduced-form within-market spillovers from rivals' supply-chain conditions (consistent with strategic complementarity and/or demand reallocation).

Results. Table 3 brings strategic interactions to the fore. As in the baseline, Columns (1)–(2) report OLS and OLS with the shift-share shortfall exposure, while the remaining columns implement the IV strategy that instruments own firm-level shortfalls with shortfall- and dwell-time exposures. Appendix Table A.6 reports the first-stage results, showing strong instruments with F-statistics comfortably above conventional thresholds.

The first three columns tell a familiar story for *own* disruptions— the IV coefficients are sizable and precisely estimated. The novel result is that *competitors'* disruptions also matter for a firm's pricing. In Column (3), the IV estimates imply that delivery delays faced by rivals pass through into a firm's prices with an elasticity around 0.12—smaller than the effect of own shocks, but economically meaningful. Firms also significantly respond to competitors' cost shocks.

Columns (4)–(5) expand the sample to include Numerator firms that never match to Panjiva, pooling importers and non-importers. Column (4) shows that the average pass-through, both own and competitors', is broadly similar to the importer-only sample. Column (5) then splits the strategic interaction terms by the firm's importing status. Even non-importing firms—those not directly exposed to supply disruptions—raise prices when their importing competitors are hit, consistent with demand reallocation toward unaffected sellers and the resulting slackening of competitive pressure.⁴⁶

These findings highlight the role of strategic interactions in shaping the broader impact of supply chain disruptions. Firms not only raised prices following their own supply shocks but also in response to competitors' disruptions, suggesting that aggregate shocks can propagate through competitive spillovers and amplify overall effects.

Next, we examine how firms' pass-through of their own and competitors' delivery delays varies with aggregate conditions. The question is whether pass-through intensifies in certain states of the economy. Table 4 addresses this by re-estimating the baseline specification from column (5) of Table 3, interacting shortfall and competitor-shortfall coefficients with indicators for different states. We

⁴⁶We cannot split the *own* shortfall and cost shocks effect by a firm's importing status because these variables are mechanically zero for non-importers.

Table 3. Pass-Through of Supply Chain Disruptions: Strategic Interactions

	Importing Firms			All Firms	
	(1) OLS	(2) OLS (Shift–Share)	(3) IV	(4) IV	(5) IV
<i>Own disruptions</i>					
Shortfall	0.006 (0.0048)	0.102*** (0.0375)	0.199*** (0.0743)	0.308*** (0.0567)	0.268*** (0.0569)
$\Delta UnitC$	0.155*** (0.0595)	0.157*** (0.0595)	0.135** (0.0608)	0.234*** (0.0550)	0.176*** (0.0586)
$\Delta FreightC$	0.102** (0.0435)	0.109** (0.0436)	0.153*** (0.0485)	0.056* (0.0328)	0.088** (0.0406)
<i>Competitor disruptions</i>					
Shortfall, compet	-0.005 (0.0129)	-0.009 (0.0126)	0.105** (0.0443)	0.124*** (0.0269)	
$\Delta UnitC$, compet	0.296* (0.1581)	0.294* (0.1580)	0.305* (0.1589)	0.374*** (0.1283)	
$\Delta FreightC$, compet	0.124 (0.1020)	0.140 (0.1022)	0.173* (0.1042)	0.057 (0.0787)	
Shortfall, compet (Imp.)					0.101*** (0.0264)
Shortfall, compet (Non-Imp.)					0.116*** (0.0303)
$\Delta UnitC$, compet (Imp.)					0.601*** (0.1435)
$\Delta UnitC$, compet (Non-Imp.)					0.018 (0.1605)
$\Delta FreightC$, compet (Imp.)					-0.008 (0.0860)
$\Delta FreightC$, compet (Non-Imp.)					0.134 (0.0871)
Firm FE	✓	✓	✓	✓	✓
Category×Quarter FE	✓	✓	✓	✓	✓
Import dummy	✓	✓	✓	✓	✓
Observations	986,023	986,023	984,692	1,705,606	1,705,606
Weak IV F-stat			406.118	750.326	760.827

Notes: The table reports regressions of 12-month price changes on own and competitor supply-chain disruptions. Columns (1)–(3) use Numerator–Panjiva matched importing firms (2020–2023). Columns (4)–(5) expand the sample to include non-importing firms. Column (1) estimates OLS using *Shortfall*. Column (2) uses shift–share *Shortfall Exposure*. Columns (3)–(5) report IV estimates instrumenting *Shortfall* with *Shortfall Exposure* and *Dwell Time Exposure*. Column (5) further splits competitor effects by importing status. All specifications include firm and product-category×quarter fixed effects. Standard errors are clustered at the product-category×quarter level. ***, **, * denote significance at the 1%, 5%, and 10% levels.

Table 4. State-Dependent Pass-Through of Delivery Shortfalls to Prices

	(1) High-Inflation Category	(2) High-Inflation Month
Shortfall $\times$ Low State	-0.688*** (0.0696)	0.262*** (0.0560)
Shortfall $\times$ High State	1.101*** (0.0670)	0.345*** (0.0591)
$\Delta UnitC$	0.181*** (0.0629)	0.218*** (0.0554)
$\Delta FreightC$	0.136*** (0.0378)	0.059* (0.0326)
Shortfall, compet $\times$ Low State	-0.162*** (0.0323)	0.105*** (0.0271)
Shortfall, compet $\times$ High State	0.345*** (0.0328)	0.138*** (0.0277)
$\Delta UnitC$, compet	0.466*** (0.1409)	0.348*** (0.1289)
$\Delta FreightC$, compet	0.141 (0.0886)	0.055 (0.0786)
Firm FE	✓	✓
Category $\times$ Quarter FE	✓	✓
Observations	1,705,606	1,705,606
Weak IV F-stat	372.299	385.171

Notes: The table reports IV regressions of 12-month product-level price changes on delivery shortfalls interacted with high- and low-inflation states. Column (1) defines the state at the product-category level (high-inflation categories versus others). Column (2) defines the state at the time level (high-inflation months versus others). All specifications instrument shortfalls using *Shortfall Exposure* and *Dwell Time Exposure*, and include current imported-input and freight cost shifters as controls. Competitor variables are constructed as leave-one-out market-share-weighted averages within product categories. Firm and product-category $\times$ quarter fixed effects are included. Standard errors are clustered at the product-category $\times$ quarter level. ***, **, * denote significance at the 1%, 5%, and 10% levels.

consider two dimensions: high-inflation sectors and periods of elevated inflation.

The first column compares high- and low-inflation sectors. A *High-inflation category* is defined as a product category-month with price growth above the annual median across all categories. Pass-through of both own and competitors' shortfalls is significantly stronger in these sectors. The second column captures periods following high inflation, when inflation perceptions and expectations are likely elevated.

A *High-inflation month* is defined as one with previous-month inflation above the annual median. Again, pass-through is higher in these periods.

Overall, the state-dependent results show that inflationary pressures amplify firms' ability to pass costs through to prices.

5 Quantification

We conduct a model-consistent accounting exercise in an impulse-response spirit. We feed the monthly paths of our supply-chain shifters into the estimated micro pricing system and trace how they propagate through within-market shares and strategic interactions into predicted consumer price changes. Because we do not observe clean structural shock series and use our shift-share exposure series—common HS- and port-level movements interacted with predetermined (2019) firm import portfolios—as our inputs, we interpret the results as an *accounting impulse response* rather than a structural decomposition of pandemic-era inflation.

To obtain decomposition, let Δp_{jt} denote the vector of equilibrium 12-month price changes for the firm-products active in market j at month t (stacked across firms). Let W_j^0 be the leave-one-out revenue-share matrix constructed from baseline (2019Q4) market shares, with off-diagonal elements $w_{ik}^0 = S_k^0 / (1 - S_i^0)$ and zeros on the diagonal. The strategic pricing system for the product market j implies⁴⁷

$$\Delta p_{jt} = \underbrace{(I - \gamma W_j^0)^{-1}}_{X_j \text{ (strategic multiplier)}} \left[\alpha \Delta mc_{jt} + \beta \Delta \tau_{jt} + \delta W_j^0 \Delta \tau_{jt} \right], \quad (8)$$

where $\gamma > 0$ governs strategic complementarities in prices, Δmc_{jt} collects marginal-cost shifters, and $\Delta \tau_{jt}$ is an *availability* index, normalized so that higher values correspond to worse availability (tighter conditions). The bracketed term in (8) is the *first-round impulse*: pass-through of own marginal costs ($\alpha \Delta mc$), own availability tightness (βA), and exposure to rivals' tightness weighted by market structure ($\delta W^0 A$). The multiplier $X_j = (I - \gamma W_j^0)^{-1} = I + \gamma W_j^0 + \gamma^2 (W_j^0)^2 + \dots$ stacks subsequent rounds of best responses.

⁴⁷A convenient way to see (8) is to start from the linearized best-response for each firm-product i in market j : $\Delta p_{ijt} = \gamma \sum_{k \neq i} w_{ik}^0 \Delta p_{kjt} + \alpha \Delta mc_{ijt} + \beta \Delta \tau_{ijt} + \delta \sum_{k \neq i} w_{ik}^0 \Delta \tau_{kjt}$. Stacking across i yields $\Delta p_{jt} = \gamma W_j^0 \Delta p_{jt} + \alpha \Delta mc_{jt} + \beta \Delta \tau_{jt} + \delta W_j^0 \Delta \tau_{jt}$, and rearranging gives (8). The share weights w_{ik}^0 correspond to the standard leave-one-out competitor-share index; see Appendix A.4 for the derivation of the weights and the system form.

To separate various channels, we label the following objects:

$$\text{Cost (direct)} = \alpha \Delta mc_{jt}, \quad \text{Cost (amplified)} = (X_j - I) \alpha \Delta mc_{jt},$$

$$\text{Availability-own (direct)} = \beta \Delta \tau_{jt}, \quad \text{Availability-own (amplified)} = (X_j - I) \beta \Delta \tau_{jt},$$

$$\text{Availability-rivals (direct)} = \delta W_j^0 \Delta \tau_{jt}, \quad \text{Availability-rivals (amplified)} = (X_j - I) \delta W_j^0 \Delta \tau_{jt}.$$

By construction, the “amplified” terms capture the additional equilibrium movement generated by strategic feedback beyond the first round.

To compute these objects, we require cost and availability shifters together with the relevant parameter values. On the cost side, we construct monthly firm-level shifters $\Delta Unit C_{ft}$ and $\Delta Freight C_{ft}$ (Section 3.3) and map them into a marginal-cost index $\Delta mc_{ft} \equiv \theta_M \Delta Unit C_{ft} + \theta_F \Delta Freight C_{ft}$. On the availability side, we set $\Delta \tau_{ft} \equiv \widehat{Shortfall}_{ft} = \hat{\pi}_S Shortfall Exposure_{ft} + \hat{\pi}_D \Delta Dwell Exposure_{ft}$, i.e. the fitted value from the first-stage relationship between realized delivery shortfalls and the two exposure measures (see Appendix Table A.6 for the weights $\hat{\pi}_S, \hat{\pi}_D$). All series are aligned by month-year, normalized to a 2019Q4 baseline, and expressed in 12-month changes.⁴⁸

We parameterize $(\alpha, \beta, \delta, \gamma, \theta_M, \theta_F)$ using our baseline micro estimates and the mapping described in Appendix A.4. We view this step as calibration/parameterization for the accounting exercise rather than structural estimation of the full model. Our parameterization reflects the baseline micro estimates in column 4 of Table 3. The interaction matrix, W_t , uses pre-pandemic revenue shares within narrowly defined product markets in Numerator.

The decomposition in Figure 8 highlights two forces—cost and availability shocks—and how strategic interactions amplify them. The fitted aggregate price series (black line) equals the sum of *cost* and *availability* components, each comprising a *direct* (first-round) effect and a *strategic* amplification from equilibrium feedback. *First*, import-related costs matter but are not the primary driver of the run-up: higher unit values and freight costs raise prices through standard pass-through, while competitors’ responses add only modest amplification. *Second*, availability tightness accounts for most of the peak. When deliveries fall short or port congestion rises, firms face tighter effective supply and increase prices; because these disruptions are widespread, price increases propagate across firms through competitive interactions, amplifying the aggregate response beyond the initial shock. *Third*, port congestion (dwell times) becomes especially important when the system is most strained, simultaneously worsening availability and strengthening amplification as many rivals are affected by the same bottleneck. *Finally*, the

⁴⁸We merge the firm-level shifters to each firm-product observation in Numerator; within a firm, all products share the same shifters in a given month, so heterogeneity in predicted price responses arises from cross-firm differences in exposure and propagation through market structure via W_j^0 and the multiplier X_j .




Figure 8. Aggregate predicted price change: costs vs. availability, own vs. strategic components

Notes: The black line plots the fitted aggregate price change from the accounting exercise. Shaded areas decompose the prediction into channels. For any month t , let $X_t = (I - \gamma W_t)^{-1}$ and $u_t = \alpha \Delta mc_t + \beta A_t + \delta W_t A_t$. Then the prediction is $\Delta p_t = X_t u_t$, where Cost Direct = $\alpha \Delta mc_t$, Cost Amp = $(X_t - I) \alpha \Delta mc_t$, Delay Own = βA_t , Amp Own = $(X_t - I) \beta A_t$, Delay Strat = $\delta W_t A_t$, and Amp Strat = $(X_t - I) \delta W_t A_t$. Firm-level shifters are merged to firm-products and aggregated using Numerator sales weights. Series are normalized to 2019Q4=0.

same mechanism generates spillovers to firms with little direct exposure: even firms that import little—or not at all—can raise prices when competing in markets where disrupted import-reliant rivals increase prices and reduce effective availability, transmitting scarcity through the competitive environment rather than solely through firms’ own exposure.

Two caveats guide interpretation. First, by treating the shortfall and dwell series as common monthly impulses, we assume proportional exposure across firms once we fix market shares. This is useful for tracing the aggregate component of supply-chain disruptions, but it abstracts from genuinely idiosyncratic, firm-specific shocks and from changes in exposure over time. Second, we hold the model objects fixed—both the pass-through parameters and the (baseline) within-market revenue shares used to construct W^0 —so we do not capture endogenous changes in pricing behavior or market structure during the episode (e.g., entry/exit, share reallocation, or state-dependent pass-through beyond what is built into the reduced-form coefficients). With these caveats in mind, we interpret the exercise as *impulse-response accounting*: it decomposes the predicted price path into (i) direct effects of cost and availability shifters and (ii) the additional amplification generated by strategic complementarities. The main take-away is that delay-driven availability tightness matters most in the run-up, and its impact is substantially magnified when many rivals face the same bottlenecks.

6 Conclusion

We study theoretically and empirically how supply chain disruptions shape consumer prices, focusing on three channels: pass-through of firms' own cost shocks, pass-through of availability shocks, and strategic responses to competitors' disruptions. On the theory side, we derive a tractable markup rule that augments the Lerner condition with an endogenous availability term and extend it to a multi-firm setting with strategic complementarities. Empirically, we build a novel large-scale dataset linking shipment-level Bills of Lading to granular consumer prices and use a shift-share instrumental-variable strategy to estimate pass-through of cost and availability shocks to retail prices. We find that both imported-input cost shocks and delivery delays raise prices materially and persistently, and that sectoral disruptions spill over across firms as they raise prices in response to competitors' shocks. An accounting exercise shows that these strategic interactions substantially amplify the price impact of supply-chain disruptions—most notably delivery delays—during the pandemic.

These findings point to supply-chain resilience as a macro-prudential concern. High-frequency indicators of inbound flows, delivery shortfalls, and port congestion could serve as early warnings when availability constraints bind, while targeted investments in port capacity, warehousing, and diversified logistics routings may help mitigate the price impact of delays. Because prices are strategic complements, congestion affecting one set of firms can raise rivals' markups and propagate price pressures across markets, underscoring the role of contestability and competition in dampening amplification. More broadly, our results also speak to inflation diagnostics for monetary policy: disruptions that operate through availability can raise effective markups and spread across firms, generating persistent inflationary pressure even when measured input-cost pass-through is modest. This suggests value in monitoring indicators of logistics tightness alongside traditional demand measures when assessing inflation persistence and the likely transmission of stabilization policy.

A natural next step is to embed our tractable pricing-inventory modeling block in a general-equilibrium framework with endogenous sourcing and inventories. Such a framework would capture how firms' collective responses—switching suppliers or transport modes and accumulating precautionary inventories—can themselves deepen congestion and amplify shocks in equilibrium. It would also allow researchers to quantify the aggregate effects of disaggregate supply-chain disruptions and evaluate policy counterfactuals such as port expansions, routing diversification, or greater competition at logistics nodes. In addition, understanding how firm size, market structure, and supply-chain diversification shape firms' ability to absorb shocks and transmit them to prices remains an important avenue for future research.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) used ChatGPT to assist with tasks such as improving clarity, editing style, and checking consistency of mathematical derivations. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

References

- Aguirregabiria, Victor.** 1999. “The Dynamics of Markups and Inventories in Retailing Firms.” *The Review of Economic Studies*, 66(2): 275–308.
- Albagli, Elías, Francesco Grigoli, Emiliano Luttini, Dagoberto Quevedo, and Marco Rojas.** 2025. “Beyond Costs: The Dominant Role of Strategic Complementarities in Pricing.”
- Alessandria, George, Joseph P. Kaboski, and Virgiliu Midrigan.** 2010. “Inventories, Lumpy Trade, and Large Devaluations.” *American Economic Review*, 100(5): 2304–2339.
- Alessandria, George, Shafaat Yar Khan, Armen Khederlarian, Carter Mix, and Kim J. Ruhl.** 2023. “The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020–2022.” *Journal of International Economics*, 146: 103788.
- Alvarez, Luis J., Emmanuel Dhyne, Marco Hoeberichts, Claudia Kwapil, Hervé Le Bihan, Patrick Lünnemann, Fernando Martins, Roberto Sabbatini, Harald Stahl, Philip Vermeulen, and Jouko Vilmunen.** 2006. “Sticky Prices in the Euro Area: A Summary of New Micro-Evidence.” *Journal of the European Economic Association*, 4(2–3): 575–584.
- Amiti, Mary, Oleg Itskhoki, and Jozef Konings.** 2019. “International Shocks, Variable Markups, and Domestic Prices.” *The Review of Economic Studies*, 86(6): 2356–2402.
- Amiti, Mary, Sebastian Heise, Fatih Karahan, and Ayşegül Şahin.** 2024. “Inflation Strikes Back: The Role of Import Competition and the Labor Market.” *NBER Macroeconomics Annual*, 38: 71–131.
- Argente, David, Salomé Baslandze, Douglas Hanley, and Sara Moreira.** 2020. “Patents to Products: Product Innovation and Firm Dynamics.” Federal Reserve Bank of Atlanta FRB Atlanta Working Paper 2020-4.
- Ascari, Guido, Dennis Bonam, and Andra Smadu.** 2024. “Global supply chain pressures, inflation, and implications for monetary policy.” *Journal of International Money and Finance*, 142: 103029.
- Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti.** 2024. “The causal effects of global supply chain disruptions on macroeconomic outcomes: Evidence and theory.” *SSRN Electron. J.*
- Baqae, David, and Emmanuel Farhi.** 2022. “Supply and demand in disaggregated Keynesian economies with an application to the COVID-19 crisis.” *Am. Econ. Rev.*, 112(5): 1397–1436.
- Baslandze, Salomé, Simon Fuchs, KC Pringle, and Michael Sparks.** 2025. “Tariffs and Consumer Prices: Insights from Newly Matched Consumption-Trade Micro Data.” *Policy Hub*, 2025(1): 1–17.

- Borusyak, Kirill, and Xavier Jaravel.** 2021. “The Distributional Effects of Trade: Theory and Evidence from the United States.” [https://www.nber.org/papers/](https://www.nber.org/papers/w28957), (28957).
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel.** 2025. “A practical guide to shift-share instruments.” *J. Econ. Perspect.*, 39(1): 181–204.
- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou.** 2024. “Investment in Infrastructure and Trade: The Case of Ports.” , (w32503).
- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou.** 2025. “Rigidities in transportation and supply chain disruptions.” *AEA Pap. Proc.*, 115: 543–550.
- Carreras-Valle, Maria-Jose, and Alessandro Ferrari.** 2025. “The Cost of Delivery Delays.” *AEA Papers and Proceedings*, 115: 618–623.
- Cavallo, Alberto, and Oleksiy Kryvtsov.** 2023. “What can stockouts tell us about inflation? Evidence from online micro data.” *J. Int. Econ.*, 103769.
- Comin, Diego A, Robert C Johnson, and Callum J Jones.** 2023. “Supply Chain Constraints and Inflation.” [https://www.nber.org/papers/](https://www.nber.org/papers/w31179), (31179).
- Davis, Steven, John Haltiwanger, and Scott Schuh.** 1998. *Job Creation and Destruction*. Vol. 1. 1 ed., The MIT Press.
- De Loecker, Jan, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik.** 2016. “Prices, markups, and trade reform.” *Econometrica*, 84(2): 445–510.
- Finck, David, Mathias Klein, and Peter Tillmann.** 2024. “The Inflationary Effects of Global Supply Chain Shocks: Evidence from Swedish Microdata.”
- Flaen, Aaron, Ali Hortaçsu, Felix Tintelnot, Nicolás Urdaneta, and Daniel Xu.** 2025. “Who pays for tariffs along the supply chain? Evidence from European wine tariffs.” National Bureau of Economic Research, Cambridge, MA:National Bureau of Economic Research.
- Flaen, Aaron, Flora Haberkorn, Logan Lewis, Anderson Monken, Justin Pierce, Rosemary Rhodes, and Madeleine Yi.** 2023. “Bill of lading data in international trade research with an application to the COVID-19 pandemic.” *Review of International Economics*, 31(3): 1146–1172.
- Franzoni, Francesco A., Mariassunta Giannetti, and Roberto Tubaldi.** 2023. “Supply Chain Shortages, Large Firms’ Market Power, and Inflation.” Swiss Finance Institute Swiss Finance Institute Research Paper Series 23-105.
- Fuchs, Simon, and Woan Foong Wong.** 2022. “Multimodal Transport Networks.” 2022(13).
- Glover, Andrew, José Mustre-del Río, and Alice von Ende-Becker.** 2023. “How much have record corporate profits contributed to recent inflation?” *Econ. Rev.*, 108(1): 1–13.
- Hacıoğlu Hoke, Sinem, Leo Feler, and Jack Chylak.** 2024. “A Better Way of Understanding the US Consumer: Decomposing Retail Sales by Household Income.”
- He, Zijian, and Yang Su.** 2023. “Measuring Welfare Gains from Online Stores:Theory and Evidence from the Supreme Court’s Wayfair Decision.” *SSRN Electronic Journal*.

- Hristakeva, Sylvia, Jura Liaukonyte, and Leo Feler.** 2024. “The no-hunger games: How GLP-1 medication adoption is changing consumer food purchases.” *Social Science Research Network*.
- Iglehart, Donald L, and Ward Whitt.** 1970. “Multiple channel queues in heavy traffic. II: sequences, networks, and batches.” *Adv. Appl. Probab.*, 2(2): 355–369.
- Khan, Aubhik, and Julia K. Thomas.** 2007. “Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies.” *American Economic Review*, 97(4): 1165–1188.
- Kingman, J F C.** 1961. “The single server queue in heavy traffic.” *Math. Proc. Camb. Philos. Soc.*, 57(4): 902–904.
- Kryvtsov, Oleksiy, and Virgiliu Midrigan.** 2013. “Inventories, Markups, and Real Rigidities in Menu Cost Models.” *Review of Economic Studies*, 80(1): 249–276.
- Lafrogne-Joussier, Raphael, Julien Martin, and Isabelle Mejean.** 2023. “Energy cost pass-through and the rise of inflation: Evidence from French manufacturing firms.” C.E.P.R. Discussion Papers CEPR Discussion Papers 18596.
- Liu, Ernest, Vladimir Smirnyagin, and Aleh Tsyvinski.** 2024. “Supply chain disruptions and supplier capital in US firms.” *Available at SSRN 4866869*.
- Lu, Patrick Alexander, Oleksiy Han, Ben Kryvtsov, and Ben Tomlin.** 2024. “Markups and inflation in oligopolistic markets: evidence from wholesale price data.”
- Manski, Charles F.** 1993. “Identification of Endogenous Social Effects: The Reflection Problem.” *Review of Economic Studies*, 60(3): 531–542.
- Meyer, Brent H., Brian C. Prescott, and Xuguang Simon Sheng.** 2023. “The impact of supply chain disruptions on business expectations during the pandemic.” *Energy Economics*, 126: 106951.
- Nakamura, Emi, and Jón Steinsson.** 2008. “Five Facts About Prices: A Reevaluation of Menu Cost Models.” *The Quarterly Journal of Economics*, 123(4): 1415–1464.
- U.S. Census Bureau.** 2026. “USA Trade Online Database.” *Foreign Trade Division*, Accessed February 2026.
- Weber, Isabella M, and Evan Wasner.** 2023. “Sellers’ inflation, profits and conflict: why can large firms hike prices in an emergency?” *Rev. Keynes. Econ.*, 11(2): 183–213.

A Additional Derivations

This appendix collects formal derivations for Section 2.1 and Section 4.2, as well as additional discussions and microfoundations.

A.1 Derivations: Simple Model (Section 2.1)

We retain the primitives and notation used in the text. Prices and marginal costs are parameterized in *logs*, p and c , with corresponding *levels* $P := e^p$ and $MC := e^c$. Demand intensity is $\lambda(p) = \Lambda e^{-\sigma p}$ with $\Lambda, \sigma > 0$. Replenishment lead times are i.i.d. exponential with rate $\mu > 0$ (mean $1/\mu$). The firm follows a base-stock policy at level $\tau \in \mathbb{N}$ (lost sales). Define

$$r \equiv \frac{\lambda}{\lambda + \mu} \in (0, 1), \quad s(\tau, \lambda, \mu) = 1 - r^\tau, \quad \kappa(\tau, \lambda, \mu) \equiv \lambda \frac{\partial s}{\partial \lambda} = -\tau(1 - r)r^\tau.$$

Accepted (served) demand arrives at rate λs . One convenient closed form for expected on-hand inventory is

$$\mathbb{E}[I(\tau, \lambda, \mu)] = \frac{\tau - r(\tau + 1) + r^{\tau+1}}{1 - r}. \quad (\text{A1})$$

Lemma 1 (Profit-rate equivalence). *Let demand be Poisson with intensity $\lambda(p) = \Lambda e^{-\sigma p}$ and lead times i.i.d. $\text{Exp}(\mu)$, independent of demand. Under a base-stock policy at level τ with lost sales, stationary/ergodic inventory, instantaneous sales at demand epochs when $I(t) > 0$, instantaneous delivery resets to τ , and a constant per-unit level margin $(P - MC)$, the long-run average contribution (revenue net of unit cost, excluding holding cost) equals margin times the throughput of accepted demand:*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \int_0^T (P - MC) \mathbf{1}\{I(t) > 0\} dN(t) = (P - MC) \lambda s(\tau, \lambda, \mu) = (P - MC) \lambda (1 - r^\tau). \quad (\text{A23})$$

Equivalently, in renewal cycles delimited by consecutive deliveries,

$$\frac{(P - MC) \mathbb{E}[\min\{N, \tau\}]}{\mathbb{E}[W]} = (P - MC) \lambda (1 - r^\tau), \quad (\text{A24})$$

where $W \sim \text{Exp}(\mu)$ is the cycle length, $N \mid W \sim \text{Pois}(\lambda W)$, and hence $N \sim \text{Geom}(\frac{\mu}{\lambda + \mu})$ on $\{0, 1, 2, \dots\}$. The full average profit rate including holding cost is

$$\Pi(p, \tau) = (P - MC) \lambda (1 - r^\tau) - h \mathbb{E}[I(\tau, \lambda, \mu)], \quad \mathbb{E}[I] \text{ given by (A1)}. \quad (\text{A25})$$

Proof of Lemma 1 (renewal-reward). Consider renewal cycles starting immediately after a delivery (on-hand = τ) and ending at the next delivery. Let $W \sim \text{Exp}(\mu)$ be the cycle length. Conditional on W , the number of demand arrivals is $N \mid W \sim \text{Pois}(\lambda W)$. Integrating out W (Poisson-exponential mixture) yields a geometric law on $\{0, 1, 2, \dots\}$ with success probability $p := \mu/(\lambda + \mu)$, so $\mathbb{P}\{N = n\} = (1 - p)^n p$ and $1 - p = r = \lambda/(\lambda + \mu)$. Sales in a cycle equal $\min\{N, \tau\}$. Using the truncated-geometric identity,

$$\mathbb{E}[\min\{N, \tau\}] = \sum_{n=0}^{\tau-1} n(1 - p)^n p + \tau(1 - p)^\tau = \frac{1 - p}{p} [1 - (1 - p)^\tau] = \frac{\lambda}{\mu} (1 - r^\tau).$$

By the renewal-reward theorem,

$$\text{avg. sales rate} = \frac{\mathbb{E}[\min\{N, \tau\}]}{\mathbb{E}[W]} = \mu \cdot \frac{\lambda}{\mu} (1 - r^\tau) = \lambda (1 - r^\tau).$$

Multiplying by $(P - MC)$ yields (A23). Subtracting the holding-cost flow $h \mathbb{E}[I]$ gives (A25). $\square^1$

Remark. The acceptance fraction used here is the *arrival-epoch* probability $s(\tau, \lambda, \mu) = 1 - r^\tau$; by PASTA it equals the time-average availability, which is the appropriate object for throughput and profit-rate calculations.

Proof of Proposition 1. Profits per unit time in *levels* are

$$\Pi(p, \tau) = (P - MC) \lambda(p) s(\tau, \lambda(p), \mu) - h \mathbb{E}[I(\tau, \lambda(p), \mu)], \quad P := e^p, \quad MC := e^c. \quad (\text{A2})$$

Step 1 (pricing FOC for given τ). Let $Q(p) := \lambda(p) s(\tau, \lambda(p), \mu)$. Then

$$\frac{\partial}{\partial p} [(P - MC) Q(p)] = P Q(p) + (P - MC) Q'(p) = 0.$$

Divide by $P Q(p) > 0$:

$$\frac{P - MC}{P} = -\frac{Q'(p)}{Q(p)} = -\frac{d \ln Q}{dp} = -\frac{d \ln \lambda}{dp} - \frac{\partial \ln s}{\partial \ln \lambda} \frac{d \ln \lambda}{dp}.$$

Since $d \ln \lambda / dp = -\sigma$ and $\partial \ln s / \partial \ln \lambda = (\lambda/s) \partial s / \partial \lambda = \kappa/s$, we obtain the generalized Lerner rule,

$$\frac{P^* - MC}{P^*} = \frac{1}{\sigma(1 + \kappa/s)}, \quad (\text{A3})$$

equivalently

$$p^* = c - \ln \left(1 - \frac{1}{\sigma(1 + \kappa/s)} \right). \quad (\text{A4})$$

Existence/uniqueness. The LHS of (A3), $1 - e^{c-p}$, is strictly increasing in p from 0 (at $p = c$) toward 1; the RHS is weakly decreasing in p because higher p lowers λ , raises s , and pushes $\kappa/s \uparrow 0$, increasing the denominator. Hence a unique solution exists.

Step 2 (discrete τ FOC). Using

$$\Delta s(\tau) = s(\tau+1) - s(\tau) = r^\tau(1-r), \quad (\text{A5})$$

and, from (A1),

$$\Delta \mathbb{E}[I] = \mathbb{E}[I]_{|\tau+1} - \mathbb{E}[I]_{|\tau} = 1 - r^{\tau+1}, \quad (\text{A6})$$

the exact discrete condition equates marginal revenue and marginal holding cost:

$$(P^* - MC) \lambda r^{\tau^*} (1-r) = h(1 - r^{\tau^*+1}). \quad (\text{A7})$$

Step 3 (continuous relaxation of τ). Let $x := r^{\tau^*} \in (0, 1)$. Using $\partial_\tau s = -x \ln r$ and $\partial_\tau \mathbb{E}[I] = 1 + \frac{r x \ln r}{1-r}$, the FOC becomes

$$(P^* - MC) \lambda (-\ln r) x = h \left(1 + \frac{r x \ln r}{1-r} \right). \quad (\text{A8})$$

This linear equation in x has the unique solution

$$r^{\tau^*} = x^* = \frac{h}{(-\ln r) \left((P^* - MC) \lambda + \frac{h r}{1-r} \right)} \in (0, 1). \quad (\text{A9})$$

¹*Alternative proof (PASTA).* Let $N(t)$ be the demand counting process (rate λ) and $L(t) := \mathbf{1}\{I(t) > 0\}$ the availability indicator. Instantaneous contribution is $(P - MC) L(t) dN(t)$. By PASTA and stationarity,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \int_0^T (P - MC) L(t) dN(t) = (P - MC) \lambda \mathbb{E}[L(0)] = (P - MC) \lambda \mathbb{P}\{I > 0\}.$$

Considering the superposed event process of demands and deliveries (i.i.d. types with $\mathbb{P}(\text{demand}) = r$, $\mathbb{P}(\text{delivery}) = 1-r$), an arrival is served iff not all of the immediately preceding τ events were demands, which occurs with probability $1 - r^\tau$. Hence $\mathbb{P}\{I > 0\} = 1 - r^\tau$ and we recover $(P - MC) \lambda (1 - r^\tau)$. Including holding cost reproduces (A25). $\square$

□

Proof of Corollary 1. We establish the availability elasticity, the sign of the price response via the implicit function theorem, the inventory response, and the limits.

Step 1 (availability elasticity). With $s = 1 - r^\tau$ and $\partial_{\ln \mu} r = -r(1 - r)$,

$$\partial_{\ln \mu} s = \frac{\partial s}{\partial r} \partial_{\ln \mu} r = (-\tau r^{\tau-1})(-r(1 - r)) = \tau r^\tau (1 - r),$$

so

$$\partial_{\ln \mu} \ln s = \frac{\tau r^\tau (1 - r)}{1 - r^\tau} > 0. \quad (\text{A11})$$

Step 2 (price response). Let $z := \kappa/s \in (-1, 0)$ and recall the levels Lerner rule (A3). Define

$$\phi(p, \mu) := \underbrace{1 - e^{c-p}}_{=(P-MC)/P} - \frac{1}{\sigma(1+z(p, \mu))} = 0. \quad (\text{A12})$$

Holding μ fixed, $\partial z / \partial p > 0$; holding p fixed, $\partial z / \partial \ln \mu > 0$. Hence

$$\phi_p = e^{c-p} + \frac{1}{\sigma} \frac{1}{(1+z)^2} \frac{\partial z}{\partial p} > 0, \quad \phi_{\ln \mu} = \frac{1}{\sigma} \frac{1}{(1+z)^2} \frac{\partial z}{\partial \ln \mu} > 0. \quad (\text{A13})$$

By the implicit function theorem,

$$\partial_{\ln \mu} p^* = -\frac{\phi_{\ln \mu}}{\phi_p} < 0. \quad (\text{A14})$$

Step 3 (inventory response). Use the discrete FOC (A7):

$$F(\tau, \mu) := (P^* - MC)\lambda r^\tau (1 - r) - h(1 - r^{\tau+1}) = 0. \quad (\text{A15})$$

Differentiate $F(\tau^*(\mu), \mu) = 0$ w.r.t. $\ln \mu$: $F_\tau \partial_{\ln \mu} \tau^* + F_{\ln \mu} = 0$, i.e.,

$$\partial_{\ln \mu} \tau^* = -\frac{F_{\ln \mu}}{F_\tau}. \quad (\text{A16})$$

We have

$$F_\tau = \ln r \left[(P^* - MC)\lambda r^\tau (1 - r) + h r^{\tau+1} \right] < 0$$

since $\ln r < 0$ and the bracket is positive; and

$$\begin{aligned} F_{\ln \mu} &= (P^* - MC)\lambda \partial_{\ln \mu} [r^\tau (1 - r)] + \lambda r^\tau (1 - r) \partial_{\ln \mu} P^* - h \partial_{\ln \mu} (1 - r^{\tau+1}) \\ &= h(1 - r^{\tau+1})[r - \tau(1 - r)] + \lambda r^\tau (1 - r) \partial_{\ln \mu} P^* - h(\tau + 1)r^{\tau+1}(1 - r) \leq 0, \end{aligned}$$

where we used $(P^* - MC)\lambda = \frac{h(1 - r^{\tau+1})}{r^\tau (1 - r)}$, $\partial_{\ln \mu} [r^\tau (1 - r)] = r^\tau (1 - r)[r - \tau(1 - r)]$, $\partial_{\ln \mu} (1 - r^{\tau+1}) = (\tau + 1)r^{\tau+1}(1 - r)$, and $\partial_{\ln \mu} P^* = P^* \partial_{\ln \mu} p^* < 0$ from (A14). Thus $F_{\ln \mu} < 0$ and, with $F_\tau < 0$, we conclude

$$\partial_{\ln \mu} \tau^* < 0. \quad (\text{A17})$$

Step 4 (limits). As $\mu \rightarrow \infty$, $r \rightarrow 0$, hence $s \rightarrow 1$ and $z \rightarrow 0^-$; (A3) implies

$$\frac{P^* - MC}{P^*} \rightarrow \frac{1}{\sigma} \Rightarrow P^* \rightarrow MC \cdot \frac{\sigma}{\sigma - 1} \quad \text{and} \quad p^* \rightarrow c + \ln \frac{\sigma}{\sigma - 1}.$$

As $\mu \rightarrow 0$, $1 - r \sim \mu/(\lambda + \mu)$, so $s = 1 - r^\tau \sim \tau(1 - r) \sim \tau \mu/(\lambda + \mu)$; moreover $\partial_{\ln \mu} z \rightarrow 0$ so the pricing response becomes locally flat, $\partial_{\ln \mu} p^* \rightarrow 0^-$. □

From generalized Lerner to a linear price rule. Start from the generalized markup condition

$$p - c = \mathcal{M}(p, \ln s) := \frac{1}{\sigma(1 + z(p, \ln s))}, \quad z := \kappa/s \in (-1, 0). \quad (D1)$$

Define $F(p, c, \ln s) := p - c - \mathcal{M}(p, \ln s) = 0$. Totally differentiating around an operating point and collecting terms gives

$$(1 - \mathcal{M}_p) dp = dc + \mathcal{M}_{\ln s} d \ln s, \quad (D2)$$

so the linear price rule is

$$dp = \alpha dc + \beta_s d \ln s + \varepsilon, \quad \alpha := \frac{1}{1 + \Gamma} \in (0, 1), \quad \beta_s := \frac{\Lambda}{1 + \Gamma} < 0, \quad (D3)$$

with $\Gamma := -\mathcal{M}_p > 0$ and $\Lambda := \mathcal{M}_{\ln s} < 0$, all evaluated at the operating point. For $\mathcal{M}(p, \ln s) = [\sigma(1 + z)]^{-1}$,

$$\mathcal{M}_p = -\frac{z_p}{\sigma(1 + z)^2}, \quad \mathcal{M}_{\ln s} = -\frac{z_{\ln s}}{\sigma(1 + z)^2} \Rightarrow \Gamma = \frac{z_p}{\sigma(1 + z)^2} > 0, \quad \Lambda = -\frac{z_{\ln s}}{\sigma(1 + z)^2} < 0. \quad (D4)$$

Interpretation: a higher price relaxes scarcity ($z_p > 0$) so own pass-through $\alpha \in (0, 1)$; higher availability raises the effective elasticity ($z_{\ln s} > 0$) so the availability coefficient $\beta_s < 0$. The specification $dp = \alpha dc + \beta_s d \ln s$ is the estimable reduced form used in the empirics.

A.2 Derivations: Strategic Interactions (Section 4.2)

A.2.1 Deriving the Markup and Stock-in Best Response

Consider the profit maximization problem of firm i written in the conjectural variation form, where the firm chooses both its log price p_i and its log stock-in shifter τ_i . The problem is given by

$$\begin{aligned} \max_{p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}} \quad & \left\{ \exp\{p_i + q_i(p_i, \mathbf{p}_{-i}, \tau_i, \boldsymbol{\tau}_{-i}; \xi) + s_i(q_i, \tau_i)\} \right. \\ & \left. - TC_i(\exp\{q_i(p_i, \mathbf{p}_{-i}, \tau_i, \boldsymbol{\tau}_{-i}; \xi)\}, \tau_i) \right\} \\ \text{subject to} \quad & h_{-i}(p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}; \xi) = 0. \end{aligned}$$

Here, p_i and q_i are the log price and log quantity demanded of firm i ; $s_i(q_i, \tau_i)$ is a function that shifts the effective stock-in probability (so that higher s_i implies improved product availability); $TC_i(\cdot, \tau_i)$ is the total cost function (in levels); and $h_{-i}(\cdot)$ is the conjectural variation vector function with elements $h_{ij}(\cdot)$ for $j \neq i$. (For brevity we omit the time subscript t .)

This formulation nests models of monopolistic competition, oligopolistic Bertrand competition, and oligopolistic Cournot competition, provided that the demand system is invertible. In particular, to capture firm behavior under monopolistic and oligopolistic Bertrand competition we may choose the conjectural variation function as

$$h_{-i}(p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}; \xi) = \begin{cases} \mathbf{p}_{-i} - \mathbf{p}_{-i}^*, \\ \boldsymbol{\tau}_{-i} - \boldsymbol{\tau}_{-i}^*, \end{cases}$$

which corresponds to the assumption that firm i believes its choices do not affect its competitors' prices (set at $\mathbf{p}_{-i}^*$) or stock-in shifters (set at $\boldsymbol{\tau}_{-i}^*$). The case of Cournot competition requires choosing $h_{-i}(\cdot)$ so that it implies $\mathbf{q}_{-i} \equiv \mathbf{q}_{-i}^*$ for some given $\mathbf{q}_{-i}^*$; with an invertible demand system, this can be ensured by setting

$$h_{-i}(p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}; \xi) = -(\mathbf{q}_{-i}(p_i, \mathbf{p}_{-i}, \tau_i, \boldsymbol{\tau}_{-i}; \xi) - \mathbf{q}_{-i}^*).$$

Thus, the firm's behavior under competition in both prices and product availability is captured by a conditional profit maximization problem.

We introduce the following notation:

1. $e^{p_i + q_i} \lambda_{ij}$ for $j \neq i$ denotes the Lagrange multipliers associated with the constraints in the maximization problem.

2. $\zeta_{ijk}(\mathbf{p}, \xi) \equiv \partial h_{ij}(\mathbf{p}, \xi) / \partial p_k$ is the elasticity of the conjectural variation function with respect to p_k , with $\zeta_{ijj}(\cdot) > 0$ as a normalization; the matrix $Z_i \equiv \{\zeta_{ijk}(\cdot)\}_{j,k \neq i}$ is assumed to have full rank. Similarly, we define

$$\theta_{ijk}(\mathbf{p}, \xi) \equiv \frac{\partial h_{ij}(\mathbf{p}, \xi)}{\partial \tau_k},$$

and denote by $\Theta_i \equiv \{\theta_{ijk}(\cdot)\}_{j,k \neq i}$ the corresponding matrix of cross-stock-in elasticities.

3. The own-price elasticity of demand is

$$\epsilon_i(\mathbf{p}, \xi) \equiv -\frac{\partial q_i(\mathbf{p}, \xi)}{\partial p_i} > 0,$$

and the cross-price elasticities are given by

$$\delta_{ij}(\mathbf{p}, \xi) \equiv \frac{\partial q_i(\mathbf{p}, \xi)}{\partial p_j} \quad (j \neq i).$$

4. Similarly, define the effective own-stock-in elasticity as

$$\eta_i(\mathbf{p}, \tau, \xi) \equiv \frac{\partial (q_i + s_i)}{\partial \tau_i} > 0,$$

and the cross-stock-in elasticities as

$$\gamma_{ij}(\mathbf{p}, \tau, \xi) \equiv \frac{\partial (q_i + s_i)}{\partial \tau_j} \quad (j \neq i).$$

The first-order conditions (after simplification) for the profit maximization problem with respect to p_i are then given by

$$\begin{aligned} (1 - \epsilon_i + \epsilon_i e^{-\mu_i}) + \sum_{k \neq i} \lambda_{ik} \zeta_{iki} &= 0, \\ \forall j \neq i \quad (-\delta_{ij} + \delta_{ij} e^{-\mu_i}) + \sum_{k \neq i} \lambda_{ik} \zeta_{ikj} &= 0, \end{aligned}$$

where the log markup is defined by $\mu_i \equiv p_i - mc_i$, with $mc_i \equiv \ln(\partial TC_i / \partial Q_i)$ being the log marginal cost. In the presence of the stock-in shifter, the effective demand depends on both q_i and s_i ; consequently, the perceived elasticity of demand is modified and given by

$$\sigma_i \equiv \epsilon_i(1 + \kappa_i) - \zeta'_i Z_i^{-1} \delta_i - \theta'_i \Theta_i^{-1} \gamma_i,$$

where $\kappa_i \equiv \partial s_i / \partial q_i$, and $\zeta_i \equiv \{\zeta_{iji}\}_{j \neq i}$ and $\delta_i \equiv \{\delta_{iji}\}_{j \neq i}$ are $(N-1) \times 1$ vectors. The matrix $Z_i \equiv \{\zeta_{ijk}\}_{j \neq i, k \neq i}$ has full rank by the invertibility assumption, and similarly for the stock-in part with $\theta_i \equiv \{\theta_{iji}\}_{j \neq i}$ and $\gamma_i \equiv \{\gamma_{ij}\}_{j \neq i}$ and the matrix $\Theta_i \equiv \{\theta_{ijk}\}_{j \neq i, k \neq i}$.

Solving the first-order conditions yields the expression for the optimal markup of firm i :

$$\mu_i = \ln \frac{\sigma_i}{\sigma_i - 1}.$$

Since ϵ_i , ζ_{ijk} , and δ_{ij} are functions of $(\mathbf{p}, \xi)$ (and now also of τ through s_i), it follows that $\sigma_i \equiv \sigma_i(\mathbf{p}, \tau; \xi)$. Therefore, we define the log markup function as

$$\mathcal{M}_i(\mathbf{p}, \tau; \xi) \equiv \ln \frac{\sigma_i(\mathbf{p}, \tau; \xi)}{\sigma_i(\mathbf{p}, \tau; \xi) - 1}.$$

Thus, the optimal price of firm i solves the fixed point equation

$$\tilde{p}_i = mc_i + \mathcal{M}_i(\tilde{p}_i, \tilde{\tau}_i, \mathbf{p}_{-i}, \tau_{-i}; \xi).$$

An analogous argument, based on the first-order condition with respect to the stock-in shifter τ_i , leads to a similar

fixed point characterization for the optimal stock-in decision. Defining the effective stock-in elasticity as

$$\phi_i \equiv \eta_i - \zeta_i^{\tau'} Z_i^{-1} \delta_i^{\tau} - \theta_i^{\tau'} \Theta_i^{-1} \gamma_i,$$

where $\eta_i \equiv \partial(q_i + s_i)/\partial \tau_i$ is the own elasticity with respect to τ_i and $\delta_i^{\tau} \equiv \{\partial q_i / \partial \tau_j\}_{j \neq i}$, and where ζ_i^{τ} and θ_i^{τ} are the corresponding derivatives of the conjectural variation function with respect to τ_i , we define the stock-in premium function by

$$\mathcal{T}_i(p_i, \tau_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}; \xi) \equiv -\ln\left(1 - \frac{mc_i^{\tau}}{\phi_i}\right),$$

with mc_i^{τ} the log marginal cost of adjusting the stock-in shifter. Hence, the optimal stock-in shifter satisfies the fixed point equation

$$\tilde{\tau}_i = mc_i^{\tau} + \mathcal{T}_i(\tilde{p}_i, \tilde{\tau}_i, \mathbf{p}_{-i}, \boldsymbol{\tau}_{-i}; \xi).$$

This completes the proof of Proposition 1 (Extended). Notice that when the stock-in shifter is absent (or when $s_i(\cdot)$ is constant), the above derivation reduces to the standard framework in which only price is chosen. In our extended model, the generalized elasticities σ_i and ϕ_i capture both the direct demand sensitivities and the strategic interdependencies across firms in the dimensions of price and product availability.

A.2.2 Deriving the Exposure Indices

Assume that product availability enters the expenditure function via a stock-in adjustment factor. In particular, let the log aggregate expenditure be given by

$$z_t = \ln \min_{\{Q_{it}\}} \left\{ \sum_{i=1}^N P_{it} Q_{it} s_{it}(\tau_{it}) \mid U(\{Q_{it}\}; Q_t) = 1 \right\}.$$

Here, $s_{it}(\tau_{it})$ adjusts the effective expenditure for product i based on its stock-in probability (or availability), with τ_{it} being the firm's stock-in shifter.

Let $E(s)$ denote the minimized expenditure:

$$E(s) = \min_{\{Q_{it}\}} \left\{ \sum_{i=1}^N P_{it} Q_{it} s_{it}(\tau_{it}) \mid U(\{Q_{it}\}; Q_t) = 1 \right\}$$

so that $z_t = \ln E(s)$. At the optimum, let Q_{it}^* be the optimal quantity for firm i . Define the Lagrangian for the cost minimization problem as

$$L(\{Q_{it}\}, \lambda; s) = \sum_{i=1}^N P_{it} Q_{it} s_{it}(\tau_{it}) - \lambda (U(\{Q_{it}\}; Q_t) - 1)$$

By the Envelope Theorem, the derivative of the minimized expenditure $E(s)$ with respect to s_{it} is given by the direct partial derivative evaluated at the optimum, i.e.,

$$\frac{\partial E(s)}{\partial s_{it}} = \frac{\partial L}{\partial s_{it}} \Big|_{Q_{it}=Q_{it}^*} = P_{it} Q_{it}^*$$

(Any indirect effects via the optimal Q_{it}^* vanish because the first-order conditions for Q_{it} ensure that $\partial L / \partial Q_{it} = 0$.) Taking the derivative of $z_t = \ln E(s)$ with respect to s_{it} , we obtain

$$\frac{\partial z_t}{\partial s_{it}} = \frac{1}{E(s)} \cdot \frac{\partial E(s)}{\partial s_{it}} = \frac{P_{it} Q_{it}^*}{\sum_{k=1}^N P_{kt} Q_{kt}^*}$$

In summary, by the Envelope Theorem applied to this cost minimization problem, differentiating z_t with respect to s_{it} gives

$$\frac{\partial z_t}{\partial s_{it}} = \frac{P_{it} Q_{it}^*}{\sum_{k=1}^N P_{kt} Q_{kt}^*} \equiv S_{it},$$

where Q_{it}^* is the optimal quantity for firm i and S_{it} is its expenditure (or revenue) market share. Now, if the availability

adjustment is given by a function $s_{it} = s(\tau_{it})$, then by the chain rule we have

$$\frac{\partial z_t}{\partial \tau_{it}} = \frac{\partial z_t}{\partial s_{it}} \cdot \frac{\partial s(\tau_{it})}{\partial \tau_{it}} = S_{it} \frac{\partial s(\tau_{it})}{\partial \tau_{it}}.$$

Next, suppose the markup function for firm i depends on competitors' stock-in shifters only through an aggregate availability index, so that we can write

$$\mathcal{M}_i(p_t, \tau_t; \xi_t) = \tilde{\mathcal{M}}_i(p_t, z(s(\tau_t)); \xi_t).$$

Then, by the chain rule the partial derivative of the markup function with respect to a competitor j 's stock-in shifter is

$$\frac{\partial \mathcal{M}_i}{\partial \tau_j} = \frac{\partial \tilde{\mathcal{M}}_i}{\partial z} \cdot \frac{\partial z}{\partial s(\tau)} \cdot \frac{\partial s(\tau)}{\partial \tau_j}.$$

By the Envelope condition we have already shown that

$$\frac{\partial z}{\partial s(\tau)} \cdot \frac{\partial s(\tau)}{\partial \tau_j} = S_{jt}.$$

Hence, it follows that

$$\frac{\partial \mathcal{M}_i}{\partial \tau_j} = \frac{\partial \tilde{\mathcal{M}}_i}{\partial z} S_{jt}.$$

When constructing the competitor delay index, we define the weight on competitor j as

$$\omega_{ijt}^\tau \equiv \frac{\frac{\partial \mathcal{M}_i}{\partial \tau_j}}{\sum_{k \neq i} \frac{\partial \mathcal{M}_i}{\partial \tau_k}} = \frac{\frac{\partial \tilde{\mathcal{M}}_i}{\partial z} S_{jt}}{\frac{\partial \tilde{\mathcal{M}}_i}{\partial z} \sum_{k \neq i} S_{kt}} = \frac{S_{jt}}{1 - S_{it}},$$

since the market shares of competitors sum to $1 - S_{it}$. Consequently, the competitor delay index is given by

$$d\tau_{-it} \equiv \sum_{j \neq i} \omega_{ijt}^\tau d\tau_{jt} = \sum_{j \neq i} \frac{S_{jt}}{1 - S_{it}} d\tau_{jt}.$$

This completes the proof: by explicitly incorporating product availability via the function $s(\tau_{it})$ into the aggregate expenditure function, we derive—using the Envelope condition—that the marginal effect of a competitor's stock-in shifter on the markup function is proportional to its market share. Therefore, the competitor delay index aggregates individual delay changes with weights $\omega_{ijt}^\tau = S_{jt}/(1 - S_{it})$, analogous to the competitor price index in the original framework.

A.3 Additional Derivations

A.3.1 Microfoundation for exponential lead times.

In the simple model we take the delivery (replenishment) lead time W to be exponential with rate μ . A standard way to rationalize both the *shape* (memorylessness) and the *scale* of W is via a heavy-traffic approximation for an upstream GI/GI/1 queue that processes replenishment orders.²

Setup. Consider an FCFS GI/GI/1 system that represents the upstream fulfillment “server.” Interarrival times have mean $1/\lambda$ and squared coefficient of variation c_a^2 ; service times S have mean $\mathbb{E}[S] = 1/\mu_s$ and squared coefficient of variation c_s^2 . Let the traffic intensity be $\rho := \lambda \mathbb{E}[S] \in (0, 1)$. For a tagged order submitted in steady state, let W denote its waiting (lead) time.

Lemma 2 (Heavy-traffic exponentiality of lead times). *As $\rho \uparrow 1$ (heavy traffic), the stationary lead time is asymptotically exponential after the canonical scaling:*

$$\frac{(1-\rho)}{\mathbb{E}[S]} W \stackrel{d}{\rightarrow} \text{Exp}\left(\frac{2}{c_a^2 + c_s^2}\right).$$

²Classical references include Kingman (1961) and Iglehart and Whitt (1970). We use the first-order stationary heavy-traffic limit and Kingman's formula as consistency checks for the mean.

Equivalently, for ρ close to one the distribution of W is well approximated by $\text{Exp}(\mu_{\text{HT}})$ with

$$\mu_{\text{HT}} \approx \frac{2(1-\rho)}{(c_a^2 + c_s^2) \mathbb{E}[S]}, \quad \mathbb{E}[W] \approx \frac{(c_a^2 + c_s^2)}{2} \frac{\mathbb{E}[S]}{1-\rho}.$$

Proof sketch. The stationary workload (and hence the waiting time under FCFS) of a GI/GI/1 queue admits a diffusion approximation in heavy traffic: the centered-scaled workload process converges to a reflected Brownian motion (RBM) with negative drift $\theta \propto (1-\rho)/\mathbb{E}[S]$ and variance parameter proportional to $c_a^2 + c_s^2$. The stationary distribution of an RBM is exponential with rate $2\theta/\sigma^2$, yielding the claimed limit for the scaled W . Matching the mean with Kingman's formula, $\mathbb{E}[W] \approx \frac{\rho}{1-\rho} \cdot \frac{c_a^2 + c_s^2}{2} \mathbb{E}[S]$, confirms the scale. $\square$

Interpretation and mapping to the simple model. Lemma 2 provides a structural justification for taking lead times to be exponential: in congested upstream environments the *entire distribution* of W is approximately memoryless, with a rate that moves one-for-one with the slack $1-\rho$ and is dampened by variability ($c_a^2 + c_s^2$). In our toy model we write $W \sim \text{Exp}(\mu)$ and treat μ as a reduced-form “replenishment speed.” The lemma implies that, under a GI/GI/1 microfoundation,

$$\mu \equiv \mu_{\text{HT}} \approx \frac{2(1-\rho)}{(c_a^2 + c_s^2) \mathbb{E}[S]},$$

so shocks that raise upstream congestion ($\rho \uparrow$) or variability (c_a^2 or $c_s^2 \uparrow$) lower μ and lengthen expected lead times $1/\mu$. From the firm's perspective—especially when it is small relative to the sector— μ is naturally taken as exogenous, summarizing the state of the upstream queue; the exponential assumption concerns the *shape*, while μ collects congestion and volatility into a single sufficient statistic.

Sanity checks. (i) In the special case of an $M/M/1$ system ($c_a^2 = c_s^2 = 1$), the heavy-traffic mean reduces to $\mathbb{E}[W] \approx \mathbb{E}[S]/(1-\rho)$, consistent with the exact $M/M/1$ formula at $\rho \uparrow 1$. (ii) Lower variability or more slack (smaller ρ) raises μ_{HT} and thus makes the exponential approximation concentrate at shorter lead times, matching the comparative statics we exploit in the toy model's pricing and inventory results.

A.3.2 Comparison with Kryvtsov and Midrigan (2013)

Pricing rules in compact form. Kryvtsov–Midrigan (KM) derive a frictionless *reset* price that multiplies a static markup by an intertemporal shadow-value term:

$$P_i(s^t) = \underbrace{\frac{\varepsilon_i(s^t)}{\varepsilon_i(s^t) - 1}}_{\text{static markup}} \times \underbrace{(1 - \delta_z) E_t \left[\frac{Q(s^{t+1})}{Q(s^t)} \Omega(s^{t+1}) \right]}_{\text{value of a marginal unit in inventory}},$$

where the *elasticity of expected sales* $\varepsilon_i(s^t)$ rises with availability (fewer stockouts), compressing markups, and the intertemporal term prices the shadow value of inventories. In our static base-stock model, availability enters *algebraically* through the levels Lerner rule (Appendix A.1, eq. (A3)):

$$\frac{P^* - MC}{P^*} = \frac{1}{\sigma(1+z)}, \quad z := \kappa/s \in (-1, 0),$$

so the effective demand elasticity is $\sigma(1+z)$. The inventory margin is summarized by the exact discrete stockout-avoidance condition (Appendix A.1, eq. (A7)):

$$(P^* - MC) \lambda r^\tau (1-r) = h(1-r^{\tau+1}),$$

which equates the marginal revenue from an extra unit of buffer to its holding cost.

Alignment and differences. First we present a self contained comparison between our model and the more fleshed out state dependent pricing model by KM:

(i) *Elasticity channel.* KM's $\varepsilon_i(s^t)$ and our $\sigma(1+z)$ capture the same mechanism: higher availability raises the effective elasticity and compresses markups. Both converge to the frictionless benchmark as stockouts vanish ($z \rightarrow 0^-$).

(ii) *Intertemporal vs. static scarcity.* KM price inventories through a discounted continuation value; in our risk-neutral steady state this reduces to the one-period stockout-avoidance margin in (A7). In both models, scarcity matters for pricing through the elasticity object that disciplines the markup.

(iii) *Units and mapping.* KM's level markup $\varepsilon/(\varepsilon - 1)$ corresponds to our gross markup $\mu = P/MC = \sigma(1 + z)/[\sigma(1 + z) - 1]$ (from (A3)). In logs, $m = p - c = -\ln(1 - 1/[\sigma(1 + z)])$; locally, both pricing rules reduce to inverse-elasticity logic.

(iv) *Scope.* KM embed inventories in a fully dynamic, nominally sticky setting where reset prices track discounted targets. We compress the same forces into closed-form static objects, delivering sharp pass-through coefficients and a linear, empirically tractable decomposition into cost and availability channels (used in the main text).

What is lost in the stationary reduction. Our closed-form environment deliberately “stationarizes” the joint pricing–inventory problem: payoffs are evaluated in a steady-state flow representation, and inventory incentives are summarized by static objects (s, κ, r) together with the stockout–avoidance margin (A7). Relative to KM's fully dynamic setup, this removes several margins:

(i) *Forward-looking inventory scarcity.* In KM, the target/reset price depends on the *expected discounted* shadow value of inventories, $E_t[\frac{Q(s^{t+1})}{Q(s^t)} \Omega(s^{t+1})]$, which varies with the aggregate state and the firm's inventory position. Our profit-rate formulation instead evaluates the shadow value at a stationary operating point, so pricing responds to *current* scarcity via the wedge $z = \kappa/s$ rather than expected future scarcity. Consequently, we do not capture anticipatory pricing/ordering when future congestion is predictable (e.g., persistent declines in μ that induce preemptive price increases and/or inventory building).

(ii) *Transition dynamics and state dependence.* KM feature a distribution of firms over inventory states and price vintages, with shocks propagating through gradual adjustment of that cross section. Our steady-state characterization abstracts from these transitions: $s(\tau, \lambda, \mu)$ is a stationary stock-in rate, not a time-path object that evolves as inventories deplete and are rebuilt. Empirically, our mapping is therefore best viewed as a local/medium-run comparative static around an operating point, not a full impulse response with lags, persistence, and state-dependent dynamics.

(iii) *Nominal rigidity and reset-price selection.* KM's pricing object is a *reset* price under nominal rigidity: only some firms reoptimize each period, and observed prices mix reset decisions with inaction. Our model keeps the posted price fixed within a replenishment cycle but otherwise collapses price-setting to a static optimum, so we omit the selection margin of which firms adjust when (and how menu costs interact with inventory states)—a mechanism that can affect measured pass-through and timing.

(iv) *Risk, discounting, and GE feedback.* The SDF ratio $Q(s^{t+1})/Q(s^t)$ in KM allows inventory values (and markups) to covary with aggregate risk and demand. Our risk-neutral, partial-equilibrium steady state suppresses these covariance terms and abstracts from general equilibrium feedback from inventories to aggregate quantities and marginal costs.

A.3.3 Deferred demand (backlogging) instead of lost sales

This subsection replaces the baseline *lost-sales* assumption with *deferred demand* (backlogging): customers who arrive when on-hand inventory is zero are not lost but join a backlog and are served later. The extension is useful for interpreting settings where stockouts mainly generate *delay* rather than foregone sales.

Replenishment opportunities and the base-stock convention. To avoid ambiguity about “exponential lead times” (which in some inventory models refers to independent order lead times with potential order crossover), we adopt the equivalent and more precise queueing statement: *replenishment opportunities arrive according to a Poisson process* with rate μ . At each replenishment opportunity the firm procures enough units to (i) clear any outstanding backlog and (ii) restore on-hand inventory to the base-stock level τ . Equivalently, the *inventory position* (on-hand minus backlog) is reset to τ at each replenishment epoch.

Let $I(t) \in \{0, 1, \dots, \tau\}$ denote on-hand inventory and $B(t) \in \mathbb{N}$ the backlog. Customer arrivals are Poisson with rate $\lambda(p)$ as in the main text.

Invariant on-hand availability. Backlogging does *not* change the law of motion for on-hand inventory $I(t)$ between replenishment epochs: $I(t)$ decrements by one upon each demand arrival until it hits zero, and then remains at zero until the next replenishment epoch resets it to τ . The backlog $B(t)$ accumulates only when $I(t) = 0$ and is cleared at the next replenishment epoch, but it does not feed back into $I(t)$ under the “clear-and-refill” convention.

Hence the stationary on-hand availability objects are *identical* to the lost-sales case. Writing

$$r \equiv \frac{\lambda}{\lambda + \mu} \in (0, 1), \quad s(\tau, \lambda, \mu) \equiv \Pr\{I > 0\} = 1 - r^\tau,$$

we retain the same demand-pressure sensitivity

$$\kappa(\tau, \lambda, \mu) \equiv \lambda \frac{\partial s}{\partial \lambda} = -\tau(1-r)r^\tau < 0.$$

Backlog in closed form. Under deferred demand, stockouts translate into waiting rather than foregone sales. By PASTA (common property of queues, i.e. Poisson Arrivals See Time Averages meaning fraction of arrivals that see the process in some state is the same as the fraction of time the process is in that state), the fraction of demand arrivals that find $I = 0$ equals $\Pr\{I = 0\} = 1 - s = r^\tau$, so the *backorder arrival rate* is $\lambda(1-s) = \lambda r^\tau$. Because replenishment opportunities arrive as a Poisson process with rate μ , a backordered customer faces an exponential residual wait with mean $1/\mu$. Little's law³ therefore yields the steady-state mean backlog

$$\mathbb{E}[B(\tau, \lambda, \mu)] = \lambda(1-s) \cdot \frac{1}{\mu} = \frac{\lambda}{\mu} r^\tau. \quad (9)$$

It is also convenient to record the comparative-static derivative with respect to the demand rate:

$$\frac{\partial}{\partial \lambda} \mathbb{E}[B(\tau, \lambda, \mu)] = \frac{1}{\mu} [(1-s) - \kappa] = \frac{r^\tau}{\mu} [1 + \tau(1-r)] > 0. \quad (10)$$

Expected on-hand inventory and marginal holding-cost savings. The expected on-hand inventory is also available in closed form under the same (r, s, κ) objects:

$$\mathbb{E}[I(\tau, \lambda, \mu)] = \sum_{i=1}^{\tau} i(1-r)r^{\tau-i} = \tau - \frac{\lambda}{\mu} s. \quad (11)$$

Differentiating with respect to λ (holding τ and μ fixed) gives

$$\frac{\partial}{\partial \lambda} \mathbb{E}[I(\tau, \lambda, \mu)] = -\frac{1}{\mu} [s + \kappa] < 0, \quad (12)$$

so higher demand depletes on-hand inventory faster, generating a marginal *savings* in holding costs. This term matters in the pricing FOC unless it is explicitly shut down.

Demand specification and elasticity (clarification). As in the main text, the demand rate is *log-linear in log price*:

$$\lambda(p) = \Lambda e^{-\sigma p}, \quad p = \ln P.$$

Equivalently, in *level* prices P this is isoelastic demand $\lambda(P) = \Lambda P^{-\sigma}$ with constant elasticity σ . This is the demand case under which the standard Lerner multiplier $\sigma/(\sigma - 1)$ applies.

Pricing with deferred demand: augmented marginal cost. Under full deferral, every arrival is eventually served, so the long-run sales flow equals $\lambda(p)$ (not $\lambda(p)s$). To capture the economic cost of delay, suppose each unit of backlog incurs a per-unit-time cost $b \geq 0$ (goodwill loss, expediting, contractual penalties, etc.), while on-hand inventory incurs holding cost $h > 0$ per unit time. The steady-state profit rate is

$$\Pi^{\text{DD}}(p, \tau) = (P - MC) \lambda(p) - h \mathbb{E}[I(\tau, \lambda(p), \mu)] - b \mathbb{E}[B(\tau, \lambda(p), \mu)].$$

Using isoelastic demand and applying the chain rule through $\lambda(p)$, the pricing FOC can be written as a standard Lerner rule against an *augmented marginal cost*:

$$P^* = \frac{\sigma}{\sigma - 1} (MC + C_\lambda), \quad C_\lambda \equiv h \frac{\partial \mathbb{E}[I]}{\partial \lambda} + b \frac{\partial \mathbb{E}[B]}{\partial \lambda}, \quad (13)$$

³A result in queueing theory stating that the average number of items in a stable system can be calculated as the arrival rate multiplied by the average time spent in the system.

where the derivatives are given in (10)–(12). Substituting those expressions yields the closed form

$$P^* = \frac{\sigma}{\sigma-1} \left(MC + \frac{b}{\mu} [(1-s)-\kappa] - \frac{h}{\mu} [s+\kappa] \right), \quad s = 1 - r^\tau, \quad \kappa = -\tau(1-r)r^\tau. \quad (14)$$

Equation (14) makes the economics transparent: slower replenishment (lower μ) raises the *shadow delay cost* term (through both $1/\mu$ and r), tending to raise desired prices, while the marginal holding-cost term is negative (inventory is lower when demand is higher), tending to lower prices.

Delay-only benchmark (tractable special case). If the goal is to isolate the *delay channel* and keep the pricing rule as close as possible to the baseline markup logic, one can explicitly abstract from marginal holding-cost savings by setting $h \partial_\lambda \mathbb{E}[I] \approx 0$ in (13). Then

$$P^* \approx \frac{\sigma}{\sigma-1} \left(MC + \frac{b}{\mu} [(1-s)-\kappa] \right),$$

so the only non-cost source of pricing pressure is the shadow backlog cost induced by congestion in replenishment.

Inventory choice under deferred demand (cost tradeoff). Under full deferral, τ affects profits only through costs (holding versus backlog). For fixed λ and μ , the discrete marginal cost of raising τ by one unit is

$$h(\mathbb{E}[I_{\tau+1}] - \mathbb{E}[I_\tau]) + b(\mathbb{E}[B_{\tau+1}] - \mathbb{E}[B_\tau]) = h - (h+b)r^{\tau+1}.$$

An interior discrete optimum therefore satisfies

$$r^{\tau^*+1} = \frac{h}{h+b}, \quad (15)$$

so higher backlog penalties ($b \uparrow$) raise the optimal buffer, while higher holding costs ($h \uparrow$) lower it.

Partial deferral bridge (lost sales nested). A simple reduced-form interpolation nests lost sales and full deferral. Suppose that when $I = 0$ an arrival is deferred with probability $\delta \in [0, 1]$ and lost with probability $1 - \delta$. Then the effective fraction of arrivals that become (eventual) sales is

$$s_\delta \equiv s + \delta(1-s) = 1 - (1-\delta)r^\tau, \quad \kappa_\delta \equiv \lambda \frac{\partial s_\delta}{\partial \lambda} = (1-\delta)\kappa.$$

Abstracting from delay costs and marginal holding-cost effects for the pricing wedge (so the only role of stockouts is truncation of sales), the generalized Lerner condition becomes

$$\frac{P^* - MC}{P^*} = \frac{1}{\sigma \left(1 + \frac{\kappa_\delta}{s_\delta} \right)} = \frac{1}{\sigma \left(1 + \frac{(1-\delta)\kappa}{\delta + (1-\delta)s} \right)}, \quad (16)$$

which nests the baseline lost-sales case at $\delta = 0$ and converges to the frictionless benchmark at $\delta \rightarrow 1$.

A note on persistence. In this tractable “clear-and-refill” deferred-demand extension, the backlog state is reset to zero at each replenishment opportunity, so there is no mechanical *inter-cycle* backlog persistence (backlog persists only within a replenishment interval). Persistence in observed prices would therefore come from persistence in the replenishment environment (e.g., shocks to μ) or from richer operational features omitted here for tractability—most notably delivery-capacity constraints or partial clearing, which allow backlog to carry over across replenishment epochs.

A.4 Quantification: model equation and implementation details

This appendix subsection documents (i) the fixed-point pricing equation used in the counterfactual simulation and (ii) the exact construction of the objects entering the simulation (shocks, competitor weights, aggregation, and the component decomposition).

Objects, indexing, and shocks. Fix a month t and a submarket s (in the code, s is the variable `combined_category`). Let $i \in \{1, \dots, n_{st}\}$ index products within the (s, t) cell. Define:

- Δp_{ist} : predicted product-level price change (stored as `dp_total` in `result`).
- Δmc_{ist} : product-level marginal-cost shock.
- $\Delta \tau_{ist}$: product-level delay/disruption shock.
- m_{ist} : product market share within submarket-month (in the code, `ms_prod`).
- ω_{st} : CPI (submarket) weight used to aggregate submarkets to the overall monthly series (in the code, `ms_category`).

In the implementation, the cost shock is constructed as a linear combination of two cost components,

$$\alpha \Delta mc_{ist} \equiv \alpha_1 dmc1_{ist} + \alpha_2 dmc2_{ist}, \quad (17)$$

where `dmc1_clean` and `dmc2_clean` are taken from `dmc1_ss` and `dmc2_ss` after missing-value handling (coalescing missing to zero) and sample-screening described below. The baseline coefficients used in the simulation are

$$\alpha_1 = 0.23, \quad \alpha_2 = 0.056.$$

and are taken from Table 3 Column (4), the reduced-form coefficients from $\Delta UnitC$ and $\Delta FreightC$ onto prices, encapsulating both the impact of unit cost and freight cost onto firm's marginal cost and the direct effect of marginal cost changes onto product-level prices.

The delay/disruption shock is constructed as

$$\Delta \tau_{ist} \equiv 0.401 \cdot dshortfall_{ss_{ist}} + 0.347 \cdot ddwell_{ist}, \quad (18)$$

again coalescing missing to zero and then applying the sample-screening rule described below. The values are taken from Table A.6, with the values reflecting the first stage coefficient of `dwell` and `shortfall` onto `shortfall`.

Behavioral equation with strategic complementarities. Within each (s, t) cell, the model assumes that product-level price changes satisfy a linear best-response equation featuring (i) direct pass-through of own cost and delay shocks, (ii) a direct “strategic delay” term that depends on competitors’ delay shocks, and (iii) strategic complementarity in price setting (a product’s price responds to competitors’ price changes). Specifically, for each product i ,

$$\Delta p_{ist} = \underbrace{\alpha \Delta mc_{ist}}_{\text{cost, direct}} + \underbrace{\beta \Delta \tau_{ist}}_{\text{delay, direct (own)}} + \underbrace{\delta \sum_{j \neq i} w_{ij,st} \Delta \tau_{jst}}_{\text{delay, direct (strategic)}} + \underbrace{\gamma \sum_{j \neq i} w_{ij,st} \Delta p_{jst}}_{\text{price complementarity}}. \quad (19)$$

The parameters used in the baseline code are

$$\beta = 0.308, \quad \gamma = 0.6, \quad \delta = \frac{0.12}{\gamma} = 0.2.$$

where the values are taken from Column (4) in 3, with β responding to the second stage coefficient of `shortfall` onto prices, γ is recovered by dividing a weighted index of the competitor index of unit cost and freight cost changes by a weighted index of their direct effect, and finally, where we recover the structural parameter δ in a theory consistent way by dividing the reduced-form effect of the competitor index of shortfalls by the estimated strategic interaction term γ .

Vector form and fixed-point solution. Stack products within a submarket-month cell into vectors:

$$\Delta p_{st} \in \mathbb{R}^{n_{st}}, \quad \Delta mc_{st} \in \mathbb{R}^{n_{st}}, \quad \Delta \tau_{st} \in \mathbb{R}^{n_{st}}.$$

Let $W_{st} \in \mathbb{R}^{n_{st} \times n_{st}}$ be the competitor-weight matrix with (i, j) entry $w_{ij,st}$ and $w_{ii,st} = 0$. Then (19) becomes

$$\Delta p_{st} = \alpha \Delta mc_{st} + \beta \Delta \tau_{st} + \delta W_{st} \Delta \tau_{st} + \gamma W_{st} \Delta p_{st}. \quad (20)$$

Rearranging yields the fixed-point system

$$(I - \gamma W_{st}) \Delta p_{st} = \alpha \Delta mc_{st} + \beta \Delta \tau_{st} + \delta W_{st} \Delta \tau_{st}, \quad (21)$$

so that

$$\Delta p_{st} = \underbrace{(I - \gamma W_{st})^{-1}}_{X_{st}} [\alpha \Delta m c_{st} + \beta \Delta \tau_{st} + \delta W_{st} \Delta \tau_{st}]. \quad (22)$$

The code constructs X_{st} as $X = \text{inv}(I - *W)$ and computes all components using matrix-vector products.

Construction of the competitor-weight matrix W_{st} . Within each (s, t) cell, the code uses product market shares m_{ist} to define competitor weights as

$$w_{ij,st} = \begin{cases} \frac{m_{jst}}{1 - m_{ist}}, & i \neq j, \\ 0, & i = j. \end{cases} \quad (23)$$

This implies each row sums to one (up to numerical tolerances): $\sum_{j \neq i} w_{ij,st} = 1$. Intuitively, $w_{ij,st}$ is the share of product j among *all competitors* of i in the same submarket-month.

For numerical stability, the denominator $1 - m_{ist}$ is floored at a small $\varepsilon = 10^{-6}$ when m_{ist} is extremely close to one (this corresponds to `eps_val = 1e-6` and the subsequent safeguard in the loop constructing W).

Component decomposition: direct vs. amplification; own vs. strategic delay. Define the multiplier $X_{st} = (I - \gamma W_{st})^{-1}$. The code reports a decomposition that separates (i) the cost channel vs. delay channel, (ii) direct effects vs. strategic amplification from price complementarities, and (iii) within delays, “own” vs. “strategic” components.

Cost channel.

$$\text{Cost Direct: } c_{st}^D \equiv \alpha \Delta m c_{st}, \quad (24)$$

$$\text{Cost Total: } c_{st}^T \equiv X_{st} (\alpha \Delta m c_{st}), \quad (25)$$

$$\text{Cost Amplification: } c_{st}^A \equiv c_{st}^T - c_{st}^D = (X_{st} - I)(\alpha \Delta m c_{st}). \quad (26)$$

Delay channel (direct: own vs. strategic).

$$\text{Delay Direct (own): } d_{st}^{D,\text{own}} \equiv \beta \Delta \tau_{st}, \quad (27)$$

$$\text{Delay Direct (strategic): } d_{st}^{D,\text{strat}} \equiv \delta W_{st} \Delta \tau_{st}, \quad (28)$$

$$\text{Delay Direct (total): } d_{st}^D \equiv d_{st}^{D,\text{own}} + d_{st}^{D,\text{strat}}. \quad (29)$$

Delay channel after price amplification (apply X_{st}).

$$\text{Delay Total (own): } d_{st}^{T,\text{own}} \equiv X_{st} (\beta \Delta \tau_{st}), \quad (30)$$

$$\text{Delay Total (strategic): } d_{st}^{T,\text{strat}} \equiv X_{st} (\delta W_{st} \Delta \tau_{st}), \quad (31)$$

$$\text{Delay Total: } d_{st}^T \equiv d_{st}^{T,\text{own}} + d_{st}^{T,\text{strat}}. \quad (32)$$

Delay amplification (own and strategic).

$$\text{Delay Amplification (own): } d_{st}^{A,\text{own}} \equiv d_{st}^{T,\text{own}} - d_{st}^{D,\text{own}} = (X_{st} - I)(\beta \Delta \tau_{st}), \quad (33)$$

$$\text{Delay Amplification (strategic): } d_{st}^{A,\text{strat}} \equiv d_{st}^{T,\text{strat}} - d_{st}^{D,\text{strat}} = (X_{st} - I)(\delta W_{st} \Delta \tau_{st}), \quad (34)$$

$$\text{Delay Amplification (total): } d_{st}^A \equiv d_{st}^T - d_{st}^D. \quad (35)$$

Finally, the total predicted product-level price change is

$$\Delta p_{st} = c_{st}^T + d_{st}^T, \quad (36)$$

which is exactly the implementation `dp_total = cost_total + delay_total`.

Sample handling and missing values. The simulation sets missing shock observations to zero using `coalesce(., 0)`. In addition, shocks are set to zero for observations outside the Panjiva matched sample:

$$\Delta m c_{ist} = \Delta \tau_{ist} = 0 \quad \text{if} \quad \text{Panjiva_match_sample}_{ist} = 0.$$

This ensures that counterfactual shocks are only applied where the underlying shipping/match data are available, while products can still interact strategically through W_{st} within a submarket-month.

Month identifiers are constructed as `date_dt = Date(year, month, 1)` (i.e., monthly frequency with day normalized to the first of the month).

Aggregation to submarket and aggregate CPI series. The code produces product-level predictions and then aggregates in two steps.

Step 1: aggregate within each submarket-month. For any product-level component z_{ist} (e.g., Δp_{ist} , c_{ist}^D , $d_{ist}^{A,own}$), the submarket-month aggregate is computed using product shares:

$$Z_{st} \equiv \sum_{i=1}^{n_{st}} m_{ist} z_{ist}. \quad (37)$$

Because m_{ist} are market shares, (37) corresponds to a weighted average when shares sum to one within (s, t) . (Earlier commented code divides by $\sum_i m_{ist}$; the active code omits the division, which is equivalent when $\sum_i m_{ist} = 1$.)

Step 2: aggregate across submarkets to an overall monthly series. Let ω_{st} denote the CPI weight of submarket s in month t . The overall monthly aggregate is

$$Z_t \equiv \frac{\sum_s \omega_{st} Z_{st}}{\sum_s \omega_{st}}. \quad (38)$$

The code applies (38) to the predicted total and each component (cost direct, cost amplification, delay direct own/strategic, delay amplification own/strategic).

B Data Appendix

B.1 Numerator Data

Numerator collects data from households in several ways. Using a mobile phone app called “Receipt Hog,” consumers can (1) snap and upload a picture of their paper receipts, (2) allow Numerator to scrape their emails for digital receipts, and (3) link loyalty and membership accounts (such as Amazon, UberEats, Walmart, and Home Depot accounts), which Numerator then scrapes for transaction information. Panelists are rewarded with coins redeemable for Amazon or Visa gift cards or for cash through PayPal. On average, Numerator rewards panelists approximately \$10 per month for providing their purchase information and completing surveys ([Hacioğlu Hoke, Feler and Chylak, 2024](#))

Our baseline sample excludes the items whose Numerator sector is either “Limited Service Restaurant,” “Restaurant,” “Non-Item,” “Unknown,” or “Indeterminate Category.” In addition, we restrict the sample to the items for which we have non-missing brand or manufacturer information and at least one non-missing product classification from *category*, *majorcat*, or *department* or a detailed product description. These restrictions give us a sample of 83,352,791 unique items sold during the period 2019-2023.⁴ For these items, we have information on their classification: sub-category, category, major category, department, and sector (listed hierarchically from more detailed to more aggregated classification). An example here would be the “Adult Cough/Chest Congestion” sub-category in the “Adult Cold Cough & Flu” category, “Cold, Cough & Flu” major category, “Personal Health Care” department, and the “Health

⁴Item ID is the finest product code available in Numerator. It is the product code that the data provider assigns based on what looks like distinct products, however because of the lack of detailed product descriptions and barcodes for many items, it is possible that exactly the same product is sold in different stores, but since it lacks barcode information, it has to be treated as different item by the Numerator. In our analysis, since we define products at the firm-brand-product category level, all these distinct item IDs that in reality describe the same product will be pooled together, getting closer to the definition of a product.

& Beauty” sector. For more than 20 million items in our sample, we have detailed product descriptions that provide more information about the packaging, weight, and other product attributes (e.g., CV MUCUS REL 30 CT). In our regression analysis, we use a benchmark classification consisting of 4,000 distinct product categories defined by the *product category* in Numerator, or the broader *major category* description if the category description is missing.


Table A.1 shows the number of stores across various distribution channels in the data, while Figure A.2 shows expenditure shares and the number of unique item IDs by broadly defined sectors in Numerator.

Table A.1. Stores: Summary Statistic of Distribution Channel

Channel	Frequency	Percent	Cum.	Channel	Frequency	Percent	Cum.
Food	10,142	20.85%	20.85%	Military	49	0.10%	98.01%
FSR - Regional/Ethnic	8,428	17.32%	38.17%	School Tuition & Fees	49	0.10%	98.11%
Gas & Convenience	7,421	15.25%	53.43%	FSR - American	46	0.09%	98.21%
Liquor	2,905	5.97%	59.40%	Sports & Recreation	46	0.09%	98.30%
Bodega	2,317	4.76%	64.16%	Mortgage Payment	44	0.09%	98.39%
Apparel	1,692	3.48%	67.64%	FSR - Midscale	41	0.08%	98.48%
Drug	1,471	3.02%	70.66%	Farm	36	0.07%	98.55%
LSR - Bakery/Cafe	1,270	2.61%	73.27%	Movie Theatre	36	0.07%	98.62%
Beauty	999	2.05%	75.33%	Outlet Store	36	0.07%	98.70%
Pet	990	2.04%	77.36%	Sports Entertainment	36	0.07%	98.77%
Home Improvement	819	1.68%	79.05%	Other Association Fees	32	0.07%	98.84%
Craft	701	1.44%	80.49%	Wireless	32	0.07%	98.90%
Online	531	1.09%	81.58%	Charities	31	0.06%	98.97%
Other	515	1.06%	82.64%	Video	31	0.06%	99.03%
LSR - Ethnic/Regional	512	1.05%	83.69%	FSR - Italian/Pizza	29	0.06%	99.09%
Dispensaries	433	0.89%	84.58%	FSR - Seafood/Steak	28	0.06%	99.15%
Sporting Goods Stores	424	0.87%	85.45%	LSR - Salad/Healthful	28	0.06%	99.21%
Postal Services	401	0.82%	86.27%	Transport Hub	28	0.06%	99.26%
Other Retail Store	390	0.80%	87.08%	Music Stores	27	0.06%	99.32%
LSR - Coffee/Bakery	385	0.79%	87.87%	Gambling	24	0.05%	99.37%
Book	357	0.73%	88.60%	Professional Services	24	0.05%	99.42%
Other Specialty Store	331	0.68%	89.28%	Veterinarians	24	0.05%	99.47%
Dollar	311	0.64%	89.92%	Public Markets	21	0.04%	99.51%
Specialty Food Retailer	300	0.62%	90.54%	Yard Services	19	0.04%	99.55%
Mass	237	0.49%	91.03%	Banks	17	0.03%	99.58%
Health	226	0.46%	91.49%	Doctors Office	17	0.03%	99.62%
Electronics	219	0.45%	91.94%	Wholesale	17	0.03%	99.65%
Shoe	203	0.42%	92.36%	Club	15	0.03%	99.69%
Baby & Toy	201	0.41%	92.77%	Parking Lot or Garage	15	0.03%	99.72%
Auto	194	0.40%	93.17%	Laundromat	13	0.03%	99.74%
LSR - Pizza	192	0.39%	93.56%	Airline	11	0.02%	99.77%
Department Store	181	0.37%	93.94%	Other Travel	11	0.02%	99.79%
FSR - Miscellaneous	181	0.37%	94.31%	Telecom	10	0.02%	99.81%
LSR - Burger	137	0.28%	94.59%	Utility Company	10	0.02%	99.83%
Discount Store	134	0.28%	94.87%	Media	9	0.02%	99.85%
Home Furnishings	128	0.26%	95.13%	Real Estate Services	9	0.02%	99.87%
Vapor Stores	121	0.25%	95.38%	Concert Hall or Theater	8	0.02%	99.88%
FSR - Sports Bar	106	0.22%	95.59%	Other Government Payments	8	0.02%	99.90%
Tobacco Shops	97	0.20%	95.79%	Child Care	7	0.01%	99.91%
Auto Services	95	0.20%	95.99%	Retail Services	7	0.01%	99.93%
Other Entertainment	89	0.18%	96.17%	Church Offerings	5	0.01%	99.94%
Other Services	85	0.17%	96.35%	Car Rental	4	0.01%	99.95%
Office	81	0.17%	96.51%	Copy Centers	4	0.01%	99.95%
Shopping Centers & Malls	74	0.15%	96.67%	Public Storage	3	0.01%	99.96%
Amusement Parks	71	0.15%	96.81%	Security Services	3	0.01%	99.97%
Spas	69	0.14%	96.95%	Taxi or Limousine	3	0.01%	99.97%
LSR - Dessert Snack	67	0.14%	97.09%	Travel Agency	3	0.01%	99.98%
Hotels & Resorts	66	0.14%	97.23%	Bakery	2	0.00%	99.98%
CloseOut	62	0.13%	97.35%	Cruise Line	2	0.00%	99.99%
Healthcare	61	0.13%	97.48%	Dance & Comedy Clubs	2	0.00%	99.99%
Florists	53	0.11%	97.59%	Bar or Night Club	1	0.00%	99.99%
LSR - Chicken	53	0.11%	97.70%	Movers	1	0.00%	100.00%
LSR - Mexican	52	0.11%	97.80%	Ticket Outlet	1	0.00%	100.00%
LSR - Sandwich/Deli	52	0.11%	97.91%	Unknown	1	0.00%	100.00%
Total	47,631	97.91%		Total	48,648	100.00%	


Notes: The table shows the number of retailers identifying with a given channel. The sample is the entire Numerator data (2018-2023) with 48,648 unique retailers and 108 unique channels. [A.1](#)

Figure A.1. Demographic Comparison between Numerator and ACS




Notes: The US Census ACS data is from IPUMS, weighted by personal weight (perwt), and covers 2019. Demographic information from Numerator data is based on a static sample of 400,000 panelists covering the period 2019-2022. For this comparison, we do not apply the Numerator demographic weights (*demo_weight*) that are designed to balance the sample further to make it representative of the US Census household characteristics.

Figure A.2. Consumer Expenditures and Products by Sectors in Numerator



Notes: (A) The figure shows the product-sector level share of total Numerator expenditures (\$4 billion from \$683 million transactions) in 2022. A rough benchmark comparison with the 2022 Consumer Expenditure Survey (CEX) suggests that the relatively large share of grocery-related spending in Numerator is expected. The CEX describes 52.6% of all expenditures as belonging to categories similar to those captured by Numerator. Among these similar categories, food comprised 69.1% of spending, while in Numerator, a similar category (Grocery, FSR (Full-Service Restaurants), LSR (Limited-Service Restaurants), Tobacco) comprised 51.4% of consumer outlays. (B) This figure shows the number of unique UPCs by product-sector from the set of all 2022 transactions from our baseline sample.

Figure A.3. Numerator (Food/Grocery) Price Change vs CPI Food


Notes: The solid red line illustrates the 12-month aggregate price changes derived from grocery and food categories in the Numerator data. We compute mean product-level price changes within each category and then aggregate these category-level changes using product category sales weights. The dashed blue line represent the Urban CPI for food items (FRED series: CPIUFDSL).

Table A.2. Top 20 and Bottom 20 Departments by Durability

Department	Sector	Durability Index	Department	Sector	Durability Index
Architecture & Design	Books	1.000	Paper & Plastic	Household	0.017
Birthday	Toys	1.000	Lottery	Non-Item	0.017
Business Office Furniture	Office	1.000	LSR Beverages	LSR	0.017
Office Lighting	Office	1.000	Packaged Bakery (Bread & Alternative)	Grocery	0.016
Posters	Books	1.000	Canned	Grocery	0.016
Sustainability	Books	1.000	Shelf Stable Meals	Grocery	0.016
Celebrate Childrens Books	Books	0.975	Alcohol Beverages	Grocery	0.014
Miscellaneous (Books)	Books	0.975	Baking & Cooking	Grocery	0.012
Parenting & Families	Books	0.966	Combustible Nicotine Products	Tobacco	0.011
Pressure & Temperature	Tools & Home Imp.	0.959	Candy (Snacks)	Grocery	0.011
Performing Arts	Books	0.945	Meat	Grocery	0.009
Romance	Books	0.911	Infant Toddler Nutrition	Baby	0.009
Carrying Cases	Office	0.910	Frozen Foods	Grocery	0.008
Biography & Memoirs	Books	0.909	Snack	Grocery	0.006
Office Furniture & Lighting	Office	0.896	Pet Food & Treats	Pet	0.006
Stationary (Baby)	Baby	0.890	Dairy	Grocery	0.005
Pregnancy & Maternity	Baby	0.859	Beverages	Grocery	0.005
Sports & Recreation (Books)	Books	0.836	Mobile App Downloads	Electronics	0.005
Travel & Nature (Books)	Books	0.833	Produce	Grocery	0.003

Notes: The table reports the 20 most and 20 least durable departments, where *durability* is defined as the inverse of the average number of transactions per user within a department. The calculation is based on users who were continuously present in the static panel for all 12 months of 2023. Tools & Home Imp: Tools and Home Improvement. LSR: Limited Service Restaurant. Tobacco: Tobacco Products and Accessories.

Figure A.4. Sectoral Heterogeneity in Price Changes



Notes: Distribution of median 12-month product-level price changes within each product category by quarter derived from the Numerator data.

B.2 BoL Panjiva Data

We construct the dataset using BoL records on U.S. imports from 2010 to 2023, taking *conname* as the importer name and aggregating its shipments to measure imports.⁵ We keep only shipments with the U.S. as the final destination⁶, retain only single-HS code shipments⁷, and drop records with redacted or missing (cleaned) firm names, using the name cleaning and standardization procedure described in the next section.

To construct moving volumes of imports, we then aggregate the data to a firm-HS2-month panel by summing shipment volumes in TEU. To ensure comparability with Numerator data, we restrict the sample to 2018–2023, balance the panel to the minimum and maximum months for each firm-HS2 pair, and drop 2018 after constructing moving cumulative measures. For each firm-HS2-month, we compute 2-, 3-, 4-, 5-, and 6-month moving sums of shipment volumes.

To convert shipment volumes to values, we merge BoL data with U.S. Trade Online data at the HS2 level, applying 2019 TEU-to-KG conversion ratios. The monthly Census data from U.S. Trade Online report containerized vessel values and weights (in kg) by 2-digit HS code. We first compute monthly unit prices for each HS2 code in 2019, along with conversion factors between weights (Census) and volumes (BoL) by dividing weight by volume for each HS2. These unit prices and conversion factors are then applied to BoL monthly import volumes at the firm-HS2 level to estimate monthly import values.

Appendix Figure A.5 compares aggregate bill-of-lading (BoL) import values with total maritime U.S. imports from the Census. The two series coincide in 2019 by construction and track each other closely thereafter, with BoL showing some underestimation in 2022.⁸ At the onset of the pandemic in early 2020, imports dropped sharply, followed by a rapid recovery and a surge in aggregate imports driven by the reopening of the economy and strong demand.

Figure A.6, in turn, further shows that much of this aggregate growth in imports was fueled by the entry of new importers. It plots total imports by firms that were already importing in 2019—before the pandemic—alongside aggregate maritime imports from the Census. Focusing only on incumbent importers, imports recover in late 2020 and 2021 but then decline over time. This contrast underscores the importance of entry and exit in the importing margin and shows that trends based on a stable sample of importers do not necessarily track aggregate import dynamics.


⁵Some importers are retailers rather than manufacturers. We verify that retailer imports are not conflated with those of third-party companies selling products through the retail chain. Inspection of the Panjiva *notifyparty* field indicates that shipments imported directly by retailers pertain to their own operations—whether for resale under private labels or for other needs. When third-party sellers ship to a retailer, the BoL typically lists the manufacturer or seller as *conname* and the retailer as *notifyparty*, often with the retailer’s name and address.

⁶Some shipments arrive at a U.S. port (e.g., Seattle) but are then forwarded to another country, such as Canada; we exclude these.

⁷Shipments with multiple HS codes cannot be split reliably across codes and account for only a small share of the sample.


⁸This likely reflects the use of 2019 conversion factors for BoL import values, whereas import costs spiked in 2022.

Figure A.5. US Imports over Time. Panjiva vs. Census



Notes: Aggregate imports in BoL Panjiva against USA Trade Online - U.S. Census Bureau imports data. Panjiva import volumes are indexed to the 2019 unit value of an HS code in Census data.

Figure A.6. US Imports over Time. Panjiva vs. Census. Panjiva Firms Importing in 2019.



Notes: Total imports of a sample of firms with positive deliveries in 2019 against USA Trade Online - U.S. Census Bureau total maritime imports data. Panjiva import volumes are indexed to the 2019 unit value of an HS code in the Census data.

B.3 Numerator-BoL Panjiva Match

We match the Numerator and BoL Panjiva datasets using company names. In BoL, we primarily use consignee names (*conname*) as the firm identifier, which often represent manufacturing companies but can also include retailers. In some cases, U.S. multinational firms ship products directly from abroad to U.S. locations, so their names may also appear as shippers in BoL. We retain the shipper names for our matching procedure, too.

In Numerator, each product entry can be associated with multiple identifiers—manufacturer, parent brand, and brand names—providing several avenues for linking to producers. Manufacturer names are often missing from scanned receipts, and the data provider prioritizes supplementing this information for high-sales products. As a result, larger firms are more likely to have manufacturer data, while it is often absent for smaller firms or less significant products. Our sample contains 9,569 unique manufacturer names, 70,205 unique parent brand names, and 85,903 unique brand names. Examples include:

- PROCTER & GAMBLE COMPANY / REGENERIST OLAY / OLAY
- PRIVATE LABEL / PRIVATE LABEL / WALMART
- LEGO SYSTEMS, INC / DISNEY LEGO / LEGO
- APPLE INC. / APPLE / APPLE WATCH SERIES 8

These manufacturer–parent brand–brand combinations provide a strong basis for matching with firm names in Panjiva. Numerator’s list of roughly 27,000 store names also helps identify Panjiva shipments received by retail stores.

We clean and standardize firm names in both Numerator and Panjiva to ensure consistent identifiers across sources, building on established routines (e.g., [Argente et al., 2020](#)) and adding steps tailored to Panjiva’s idiosyncrasies. The cleaning begins with basic normalization: we convert names to uppercase, remove punctuation and special characters, normalize variants of “AND” (e.g., AND, ET, UND) to a common token, trim extra whitespace, and harmonize common abbreviations. We then strip legal entity designators and other corporate suffixes (e.g., INC, LLC, LTD, CORP, GMBH, S.A.) and handle common misspellings observed in our datasets (e.g., ENTERPRIS → ENTERPRISE, APAREL → APPAREL, INCCORPORATED → INCORPORATED). For Panjiva, we additionally remove address-like information that is frequently embedded in consignee strings. Using the Panjiva location fields (city, state, postal code, country), we sequentially drop trailing location tokens (country → ZIP → state → city), as well as PO BOX strings and other distribution/warehouse phrases that often appear after the firm name. For example, TARGET ATLANTA GEORGIA 30309 USA is reduced to TARGET, and LG ELECTRONICS PANAMA S.A. is reduced to LG ELECTRONICS. Finally, we construct a no-spaces (“glued”) version of each standardized name to improve robustness to spacing-related variation during matching.

Matching between Numerator and Panjiva begins with the Numerator data, where each observation includes the manufacturer, parent brand, and brand names. We perform exact name matching with standardized Panjiva consignee and shipper names using a hierarchical approach: first matching on manufacturer names, then parent brands, and finally brand names to consignee names. If no match is found, we repeat the process using shipper names. This procedure yields a firm-level linkage between the two datasets, enabling the analysis of firm-level price changes in relation to import activity. The matched Numerator–Panjiva sample relies on this linkage and identifies the set of Numerator firms that are importers. For non-importing (unmatched) firms in Numerator, we continue to use manufacturer and parent brand names to define firm identifiers, but drop observations where only the brand name is available to reduce noise.

Table A.3. Summary Statistics at the Product-Month Level. Importing vs. Non-Importing Firms


	Numerator-Panjiva Sample (Importing Firms)	Unmatched Numerator Sample (Non-Importing Firms)
Number of Observations	1,542,111	863,795
Date Range	2019 - 2023	2019 - 2023
Distinct Number of Products	50,375	40,000
Distinct Number of Firms	4,174	17,780
Distinct Number of Product-Categories	3,633	3,381
Distinct Number of Product-Departments	213	203
Average Price Changes	0.041	0.035
Average Sales	1,713,221	659,752
Average Number of Transactions (All Users)	915	313
Average Number of Transactions (Static Users)	439	163
Panjiva Firm Dummy	1.000	0.000

Notes: This table presents summary statistics at the product-month level from 2019 to 2023. Column (1) shows statistics for the baseline Numerator-Panjiva matched sample (with Panjiva cleaning procedures as described in Section 3.2). Column (2) shows statistics for the unmatched Numerator Firms (after cleaning procedures as described in Section 3.1).


C Additional Empirical Results

Figure A.7. Variation in Shifts and Shares

(A) Aggregate HS code shortfalls over time




(B) Firm heterogeneity in HS-code sourcing within product categories




Notes: (A) Delivery shortfalls constructed as mean level deviations in the same month compared to the 2019 baseline, HS2 categories weighted with firm-level value shares. (B) The figure shows heterogeneity in firms' reliance on imported inputs within narrowly defined product categories. For each HS2 code and product category, we compute the standard deviation of firms' 2019 import shares across HS2 codes. We then average this dispersion across categories for each HS2 code. The y-axis therefore reflects how differently firms within the same product category rely on a given HS2 input. The sample is restricted to category-HS2 cells with at least 10 firms.

Figure A.8. Price Changes, Firm Size, and Supply Chain Diversification


(A) Product Price Growth vs Firm Size




(B) Product Price Growth vs Firm Market Share



(C) Product Price Growth vs Supplier Countries



(D) Product Price Growth vs Product(HS2)-Supplier Countries



Notes: (A) Binscatters of product-level price changes against log firm size, defined as the average yearly sales of the firm in Numerator, residualized for year and department fixed effects. (B) Binscatters of product-level price changes against firm market share, defined as the firm's average sales share in its product category in a year, residualized for year and department fixed effects. (C) Binscatters of product-level price changes against the log number of countries the firm sources from in a year, controlling for firm size, year, and department fixed effects. (D) Binscatters of product-level price changes against the log average number of countries the firm sources from for various product codes (HS2) in a year, controlling for firm size, year, and department fixed effects.

Table A.4. Price Effects of Own Supply Chain Disruptions: I-stage Estimates

	Col(3) I-stage Shortfall (1)	Col(4) I-stage Shortfall (2)	Lag Shortfall (3)	Col(5) I-stage Shortfall (4)	Lag Shortfall (5)
Shortfall exposure	0.491*** (0.0177)	0.470*** (0.0161)	-0.038** (0.0164)	0.425*** (0.0160)	-0.070*** (0.0164)
$\Delta Dwell$	0.168*** (0.0143)	0.151*** (0.0144)	0.056*** (0.0155)	0.140*** (0.0142)	0.048*** (0.0155)
$\Delta UnitC$	0.165*** (0.0262)	0.093*** (0.0254)	0.056** (0.0251)	0.111*** (0.0252)	0.069*** (0.0250)
$\Delta FreightC$	-0.223*** (0.0219)	-0.218*** (0.0231)	-0.005 (0.0212)	-0.228*** (0.0228)	-0.011 (0.0212)
Lag Shortfall exposure		0.070*** (0.0159)	0.531*** (0.0168)	0.088*** (0.0157)	0.543*** (0.0167)
Lag $\Delta Dwell$		0.015 (0.0126)	0.109*** (0.0135)	0.024** (0.0124)	0.116*** (0.0135)
Lag $\Delta UnitC$		0.138*** (0.0242)	0.173*** (0.0257)	0.116*** (0.0240)	0.157*** (0.0256)
Lag $\Delta FreightC$		-0.016 (0.0230)	-0.146*** (0.0232)	-0.016 (0.0228)	-0.146*** (0.0231)
Observations	991,619	962,570	962,570	962,570	962,570
F-test of excl. instruments	395.15	216.07	276.55	181.89	297.83

Notes: The table reports I-stage estimates for columns (3)-(5) in Table 2. Standard errors, clustered at the product category–quarter level, are reported in parentheses. ***, **, *: significance at the 1%, 5%, 10% levels, respectively.

Table A.5. Price Effects of Own Supply Chain Disruptions: By Product Durability

	Δp		
	(Benchmark IV)	(Durability Index)	(Durability Dummy)
	(1)	(2)	(3)
Shortfall	0.253*** (0.0802)	0.106 (0.1020)	0.234*** (0.0782)
Shortfall $\times$ Durability Index		3.445* (2.0303)	
Shortfall $\times$ Durability Dummy			0.102 (0.0958)
$\Delta UnitC$	0.018 (0.0659)	0.001 (0.0672)	0.008 (0.0665)
$\Delta FreightC$	0.027 (0.0578)	-0.002 (0.0590)	0.023 (0.0577)
Lag $\Delta UnitC$	0.196*** (0.0657)	0.188*** (0.0664)	0.188*** (0.0662)
Lag $\Delta FreightC$	0.230*** (0.0543)	0.230*** (0.0544)	0.232*** (0.0543)
Firm FE	✓	✓	✓
Cat-Quarter FE	✓	✓	✓
Observations	962570	962570	962570
Weak IV F-stat	344.742	23.147	196.335

Notes: The table reports regressions analogous to the benchmark IV own-pass-through estimate in column (5) of Table 2, but with shortfalls interacted with the *Durability Index* (column 2) and the *Durability Dummy* (column 3). Product durability is constructed based on the product purchase frequency, following Argente et al. (2020). Standard errors, clustered at the product category–quarter level, are reported in parentheses. ***, **, *: significance at the 1%, 5%, 10% levels, respectively.

Table A.6. Strategic Interactions: I-Stage Estimates

	Col(3) I-stage Shortfall (1)	Col(4) I-stage Shortfall (2)	Col(5) I-stage Shortfall (3)
Shortfall exposure	0.486*** (0.0175)	0.401*** (0.0156)	0.393*** (0.0156)
$\Delta Dwell$	0.181*** (0.0137)	0.347*** (0.0101)	0.361*** (0.0103)
$\Delta UnitC$	0.147*** (0.0249)	0.110*** (0.0235)	0.149*** (0.0244)
$\Delta FreightC$	-0.242*** (0.0207)	-0.376*** (0.0109)	-0.297*** (0.0166)
Shortfall, compet	-0.575*** (0.0078)	-0.442*** (0.0062)	
$\Delta UnitC$, compet	-0.021 (0.0641)	0.028 (0.0477)	
$\Delta FreightC$, compet	-0.191*** (0.0457)	-0.141*** (0.0327)	
Shortfall, compet Imp.			-0.424*** (0.0067)
Shortfall, compet Non-Imp.			-0.469*** (0.0067)
$\Delta UnitC$, compet Imp.			-0.103* (0.0562)
$\Delta UnitC$, compet Non-Imp.			0.242*** (0.0488)
$\Delta FreightC$, compet Imp.			-0.232*** (0.0365)
$\Delta FreightC$, compet Non-Imp.			-0.033 (0.0344)
Observations	984692	1705606	1705606
F-test of excl. instruments	406.12	750.33	760.83

Notes: The table reports I-stage estimates for columns (3)-(5) in Table 3. Standard errors, clustered at the product category–quarter level, are reported in parentheses. ***, **, *: significance at the 1%, 5%, 10% levels, respectively.

Table A.7. Pass-Through of Supply Chain Disruptions: More Detailed Import Codes

	<i>Importing firms</i>					<i>All firms</i>	
	(1) (OLS-Shift Share)	(2) (IV)	(3) (IV)	(4) (OLS-Shift Share)	(5) (IV)	(6) (IV)	(7) (IV)
Shortfall exposure	0.114*** (0.0370)	0.296*** (0.1056)	0.315*** (0.1129)	0.116*** (0.0372)	0.305*** (0.1066)	0.352*** (0.0628)	0.326*** (0.0644)
$\Delta UnitC$	0.092** (0.0417)	0.071* (0.0427)	0.070 (0.0428)	0.102** (0.0422)	0.090** (0.0428)	0.147*** (0.0405)	0.126*** (0.0410)
$\Delta FreightC$	0.124*** (0.0406)	0.180*** (0.0495)	0.183*** (0.0503)	0.126*** (0.0407)	0.184*** (0.0497)	0.064* (0.0327)	0.107*** (0.0401)
Shortfall, compet				-0.007 (0.0146)	0.169*** (0.0630)	0.154*** (0.0304)	
$\Delta UnitC$, compet				0.140 (0.1156)	0.186 (0.1167)	0.158* (0.0950)	
$\Delta FreightC$, compet				0.133 (0.1006)	0.177* (0.1043)	0.036 (0.0788)	
Shortfall, compet Imp.							0.129*** (0.0299)
Shortfall, compet Non-Imp.							0.167*** (0.0355)
$\Delta UnitC$, compet Imp.							0.324*** (0.1058)
$\Delta UnitC$, compet Non-Imp.							-0.168 (0.1260)
$\Delta FreightC$, compet Imp.							-0.032 (0.0853)
$\Delta FreightC$, compet Non-Imp.							0.130 (0.0875)
Firm FE	✓	✓	✓	✓	✓	✓	✓
Cat-Quarter FE	✓	✓	✓	✓	✓	✓	✓
Import dummy			✓	✓	✓	✓	✓
Observations	1006089	995719	995719	999110	988762	1842683	1842683
Weak IV F-stat		256.109	226.951		255.719	780.740	783.508

Notes: The table reports regressions of 12-month price changes on measures of own and competitors' supply-chain disruptions, estimated using product-month-level data. Columns (1)–(5) use the sample of importing firms (Numerator–Panjiva matched data) for 2020–2023, while columns (6)–(7) additionally include non-importing firms that never match to Panjiva. Columns (1) and (4) estimate OLS specifications using the *Shortfall Exposure* measure; the remaining columns report IV estimates that use *Shortfall Exposure* and *Dwell Time Exposure* as instruments for *Shortfall*. Firm-level shortfall, unit-cost, and freight-cost exposure variables are constructed using HS4-code aggregation; to mitigate HS4-level classification noise in Panjiva, shortfall exposure is based on HS2-level shifts. All specifications include an import dummy, firm, and product-category–quarter fixed effects. Standard errors, clustered at the product-category–quarter level, are reported in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.