

Shipping Cost Shocks and US Producer Prices: Normal versus Crisis Pass-through

Simon Gilchrist, Bin Wei, Vivian Z. Yue, and Egon Zakrajšek

Working Paper 2026-13
September 2026

Abstract: This paper examines how shipping costs pass through to US producer prices. We link confidential shipment-level import transactions from the Census Bureau’s Longitudinal Firm Trade Transactions Database (LFTTD) to confidential firm-level producer prices from the US Bureau of Labor Statistics’ Producer Price Index (PPI) microdata, constructing a monthly panel of more than 6,000 US manufacturing importers from 2005 to 2022. Using local projections and Bartik-style instruments based on route-level shipping-cost variation, we find that pass-through to producer prices is negligible in normal times but rises sharply during the COVID-19 pandemic, when shipping costs surged broadly across routes and source countries. We interpret this state dependence through a menu-cost model with imported intermediate inputs and endogenous substitution between imported and domestic inputs.

JEL classification: E31, F14, F41, L11

Key words: supply chain disruption, shipping costs, inflation, pass-through, producer prices

<https://doi.org/10.29338/wp2026-13>

The views expressed in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Federal Reserve Banks of Atlanta and Boston, the Federal Reserve System, the US Census Bureau, or the Bureau of Labor Statistics (BLS). The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was conducted with restricted access to BLS data and was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2878 (CBDRB-FY25-0470). Prepared for the Carnegie-Rochester-NYU Conference Series on Public Policy, New York, April 30–May 1, 2027.

Simon Gilchrist (sg40@nyu.edu) is with New York University’s Department of Economics and the National Bureau of Economic Research. Bin Wei (bin.wei@atl.frb.org) is in the Research Department of the Federal Reserve Bank of Atlanta. Vivian Z. Yue (vyue@emory.edu) is with Emory University’s Department of Economics, the National Bureau of Economic Research, and the Center for Economic Policy and Research. Egon Zakrajšek (egon.zakrajsek@bos.frb.org) is with the Federal Reserve Bank of Boston and the Center for Economic Policy and Research.

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/subscribe.

1 Introduction

Understanding how trade costs transmit to domestic prices is a first-order question for macroeconomics and trade policy. Shipping costs, tariffs, and other border-related charges constitute a significant component of the total cost of imported intermediate inputs for U.S. manufacturers. Yet despite the importance of this channel, the empirical literature has struggled to measure the degree to which these costs are passed through to the prices that domestic producers charge for their output. The difficulty arises, in part, from a lack of datasets that simultaneously capture firm-level trade costs and firm-level output prices at sufficient granularity.

In this paper, we address this gap by constructing a novel panel dataset that links two confidential microdata sources: the Longitudinal Firm Trade Transactions Database (LFTTD) from the U.S. Census Bureau, which provides shipment-level information on shipping costs and duties for individual importing firms, and the Producer Price Index (PPI) microdata from the Bureau of Labor Statistics (BLS), which provides firm-level output prices. By matching these databases using a multi-step name-matching algorithm, we obtain a panel of more than 6,000 U.S. manufacturing importing firms observed at a monthly frequency between 2005 and 2022, yielding roughly one million firm-month observations. To the best of our knowledge, this is the first study that directly links firm-level producer prices to detailed, transaction-level trade cost measures.

Our empirical analysis yields two main findings. First, during normal times, the pass-through of shipping costs to producer prices is negligible. OLS regressions of future log producer prices on current log ad valorem shipping costs—controlling for industry-by-month fixed effects, firm fixed effects, and lagged prices—yield coefficients that are economically small and statistically insignificant. This finding suggests that, in

normal times, fluctuations in shipping costs are largely absorbed before they translate into output prices. One interpretation, formalized in our model, is that idiosyncratic shipping-cost movements largely average out across firms’ diversified route portfolios, and firms can partly offset what remains by adjusting the mix of imported and domestic inputs, so that little of the movement ever reaches marginal cost.

Second, during the COVID-19 pandemic, when global supply chains experienced unprecedented disruption and shipping costs surged, the pass-through of shipping costs to producer prices increased dramatically. Our preferred instrumental variables (IV) estimates, which use Bartik-style instruments constructed from firms’ predetermined exposure to trade routes interacted with leave-one-out route-level shipping costs, indicate that the pass-through during the pandemic was an order of magnitude larger than in normal times. In Section 4 we use the model to translate these estimates into an aggregate inflation contribution.

The differential pass-through across episodes—negligible in normal times, amplified during the pandemic—reflects the nature of the underlying shocks and the adjustment margins available to firms. During the pandemic, shipping cost increases were broad-based and highly correlated across trade routes and source countries. The combination of large cost increases and diminished flexibility generated substantially higher pass-through to output prices.

Our identification strategy addresses a key endogeneity concern: shipping costs may be correlated with demand conditions that independently affect producer prices. To isolate exogenous supply-side variation in trade costs, we construct Bartik-style instruments that interact each firm’s historical (pre-determined) trade route shares with national average shipping costs on each route, excluding firms in the same 2-digit NAICS industry (“leave-one-out”). The IV estimates are substantially larger than the corresponding OLS estimates, consistent with downward bias in OLS arising from

demand shocks that simultaneously reduce both import volumes (and hence measured shipping costs) and producer prices.

To interpret these empirical findings, we develop a menu-cost model of shipping-cost pass-through with imported intermediate inputs and endogenous substitution between imported and domestic inputs. The model is designed to rationalize three facts that are difficult to reconcile in standard frameworks: shipping-cost pass-through is negligible in normal times; it rises sharply during the pandemic; and the estimated crisis pass-through increases with the horizon, reaching its largest value after twelve months. Our framework combines price adjustment frictions with an import-intensity margin, allowing firms to adjust the foreign–domestic composition of inputs while keeping the set of imported varieties and routes fixed. A distinguishing feature of our approach is that every pass-through result is derived in closed form.

Our model hinges on the following central mechanism. In normal times, shipping-cost movements are largely idiosyncratic across routes and source countries, and transitory. Diversification does most of the work: fluctuations that are idiosyncratic across routes largely cancel within the firm’s route portfolio, so only a small residual enters its import costs. Input substitution—the firm shifting toward domestic inputs wherever delivered import prices rise—trims that residual further, though, as we show below, its quantitative contribution is modest. Almost nothing reaches marginal cost, and what never reaches marginal cost cannot be passed through to prices, matching the insignificant estimates outside the crisis.

The pandemic shock is different in both dimensions. We model it as an MIT shock: an initially unanticipated, deterministic path for the *common* component of shipping costs, perfectly foreseen once it arrives. Because shipping costs rise simultaneously across routes, diversification offers no protection, and input substitution offsets the shock only partially: the common component passes into marginal cost with an elastic-

ity equal to the product of the intermediate share and the import share. Our central theoretical result characterizes the resulting crisis pass-through in closed form. The effect of the current shipping shock on a firm’s expected price h months ahead equals the fraction of the shock that prices will have absorbed by then, per unit of the current shock. This absorption fraction has two ingredients: which firms will have repriced by $t + h$ — repricing opportunities arrive at the calibrated monthly frequency—and how much of the shock each repricing builds in, namely its expected remaining path rather than its current value, so that anticipated increases enter prices before they occur.

We discipline the model with the same variation that underlies the empirical estimates. The elasticity of substitution between domestic and foreign inputs is recovered from the first-stage IV relationship; the menu-cost parameters are calibrated to match the pre-pandemic frequency and average absolute size of producer-price changes in the PPI microdata. We then evaluate the closed form for a benchmark jump–decay shock—a jump on impact followed by geometric decay at a constant monthly rate. Two findings emerge. First, persistence delivers the empirical sign and shape: when the shock outlives the average price spell, incorporation is still building at the horizons we estimate, so the pass-through profile is positive and *rising* with the horizon, as in the IV estimates. Second, the calibrated magnitudes fall short of the IV profile by a factor of three to five, and the closed form implies a ceiling—the static delivered-cost loading of measured freight—that the 12-month IV estimate exceeds by a factor of four.

The magnitude gap is informative. Within the model, no persistence of the measured freight shock can rationalize the size of the estimated crisis coefficients: measured freight and insurance charges must understate the effective delivered-cost disruption of the pandemic. Port congestion, delivery delays, expediting, and inventory costs rose together with freight rates, and the estimated coefficients are consistent with prices re-

sponding to this broader disruption, for which measured freight is the observable proxy. The model emphasizes the economics behind the empirical estimates. The pandemic mattered not only because shipping costs increased, but because they increased in a broad-based, persistent way.

Related Literature

Our paper connects to several strands of the literature in international trade, macroeconomics, and industrial organization. We organize our discussion around three broad themes: the pass-through of trade costs to prices, the macroeconomic effects of supply chain disruptions, and firm-level pricing behavior.

Trade Cost Pass-through. A large literature in international economics studies how changes in border costs transmit to domestic prices. Much of this work focuses on exchange-rate movements, and a central finding is that pass-through is typically incomplete: changes in costs at the border are only partially reflected in downstream prices, with the degree of pass-through varying across industries and declining along the pricing chain from import to retail prices ([Goldberg and Knetter, 1997](#)). Subsequent work documents substantial heterogeneity in pass-through across countries ([Campa and Goldberg, 2005](#)) and shows that invoicing currency is central under dominant-currency pricing ([Gopinath et al., 2010](#)). More recently, [Cuba-Borda et al. \(2025\)](#) study the inflationary effects of trade-cost shocks and show that increases in trade costs raise consumer-price inflation, with shocks to intermediate-input trade costs generating particularly persistent effects.

The pass-through of shipping costs to domestic producer prices has received less attention, in part because it requires data that jointly measure transportation costs and prices at high frequency and at the firm level. Existing work has studied freight costs

using detailed trade data and has shown that transportation costs are an important component of trade costs and trade patterns (Hummels, 2007). Most existing evidence on the inflationary effects of shipping costs is based on aggregate price indexes rather than prices at the firm level. Carrière-Swallow et al. (2023) use cross-country time-series data to study the response of aggregate import-price, producer-price, and consumer-price indexes to global shipping-cost shocks. Michail et al. (2022) similarly examine the pass-through of freight rates to aggregate and category-level consumer price indexes in the euro area. Related work examines freight costs and import prices (Hummels et al., 2009) and shipping costs and trade volumes (Shapiro, 2016). In contrast, we link realized shipment-level shipping costs directly to the producer prices charged by the same importing firms. Our contribution is to link shipment-level measures of shipping costs directly to firm-level producer prices. This allows us to estimate how shocks to the cost of imported intermediate inputs pass through to U.S. manufacturing prices at an early stage of the pricing chain.

Supply Chain Disruptions and Inflation. The COVID-19 pandemic generated intense interest in the macroeconomic effects of supply chain disruptions, and a growing literature documents that the pandemic-era disruptions—port congestion, container shortages, and shipping delays—contributed to the global surge in inflation observed in 2021–2022. At the aggregate level, Di Giovanni et al. (2022) develop a multi-country, multi-sector New Keynesian framework to decompose the drivers of pandemic inflation and find that supply-side factors, including supply chain bottlenecks, account for a substantial share of the surge, with effects propagating across countries through international production linkages. Bonadio et al. (2021) use a related multi-country, multi-sector model to quantify the global output losses from pandemic-related supply chain disruptions, attributing roughly a quarter of the model-implied GDP decline to

transmission through cross-border supply chains. A complementary literature studies the transmission of trade costs to prices. [Amiti et al. \(2021\)](#) document that the pass-through of import prices to U.S. domestic producer prices rose markedly after the onset of the pandemic, a given change in import prices generating a larger response of the producer price index than in the pre-pandemic period. At the firm level, [Lafrogne-Joussier et al. \(2023\)](#) use the micro-data underlying the French producer price index to show that firms more exposed to imported-input and energy cost shocks raised their output prices by more, estimating that a 10 percent increase in foreign costs raises output prices by roughly 0.74 percent. A growing micro-level literature studies the consumer price effects of supply-chain disruptions through delivery delays and product availability. [Cavallo and Kryvtsov \(2023\)](#) use online retail data to show that unexpected stockouts raise consumer-price inflation, with larger and more persistent effects for imported goods and import-intensive sectors. Most closely related, [Baslandze and Fuchs \(2026\)](#) link bill-of-lading import records to granular consumer transaction prices and find that both import-cost shocks and delivery shortfalls raise retail prices. [Jiao et al. \(2026\)](#), using AIS-based measures of port-to-port shipping times, likewise find that shipping delays raise consumer prices. These studies focus on consumer prices and the availability consequences of supply-chain disruptions, whereas we directly measure firm-level shipping costs and study their state-dependent pass-through to U.S. producer prices.

Our paper complements this literature by providing firm-level evidence on the mechanism through which shipping costs transmit to producer prices in the United States. By exploiting within-firm variation in shipping costs, we distinguish the pass-through from compositional changes in the types of firms that face higher trade costs. Our finding that the pass-through was dramatically amplified during the pandemic—but negligible in normal times—highlights the state-dependent nature of cost pass-through

and suggests that the inflationary effects of trade-cost shocks depend critically on the breadth and severity of the underlying disruption.

Firm-Level Pricing. Price stickiness shapes the dynamics of pass-through. [Nakamura and Steinsson \(2008\)](#) document substantial heterogeneity in the frequency of price adjustment, a natural source of the delayed, horizon-dependent pass-through we estimate, and [Gilchrist et al. \(2017\)](#) use the same PPI microdata we employ to show that financially distressed firms raised prices during the Great Recession. A state-dependent literature stresses that which and how many firms adjust matters for aggregate pass-through ([Midrigan, 2011](#)): large common shocks raise inflation partly by inducing more firms to adjust, generating a nonlinear, state-dependent relationship between costs and prices ([Blanco et al., 2024b](#)). This frequency channel can be self-reinforcing—higher inflation raises the fraction of price changes, which feeds back into inflation—steepening the cost–price relationship precisely when shocks are large ([Blanco et al., 2024a](#)). Our PPI evidence fits this mechanism—broad shipping-cost shocks raise both the frequency and the size of upward price changes, generating larger pass-through than idiosyncratic shocks that firms buffer through sourcing substitution.

Most closely related is [Gagliardone et al. \(2025\)](#), who construct an empirical measure of firm-level price gaps and show that, through a generalized adjustment hazard, these gaps discipline the frequency and size of price changes; state dependence matters little for small shocks but becomes decisive during inflation surges, when aggregate shocks push the gap distribution into the steep region of the hazard. We adopt this framework but apply it to an externally measurable cost source—firm-level exposure to shipping shocks—augmenting it with imported intermediates and endogenous sourcing. This lets us separate idiosyncratic trade-cost shocks, which firms buffer through supplier substitution and markups, from common shipping shocks, which shift the gap

distribution systematically and generate large crisis-period pass-through. Our contribution is to identify how global supply-chain disruptions enter that pricing framework and transmit to U.S. producer prices.

The remainder of the paper is organized as follows. Section 2 describes our data sources, measurement approach, and empirical results. Section 3 outlines a theoretical framework that rationalizes our empirical findings. Section 4 calibrates the model and evaluates the closed-form crisis pass-through against the IV estimates. Section 5 concludes.

2 Empirical Analysis

2.1 Data and Measurement

The key novel contribution of our paper is to construct firm-level trade cost and price inflation measures by combining microdata from two confidential databases maintained by the U.S. Census Bureau and the Bureau of Labor Statistics.

Longitudinal Foreign Trade Transaction Database (LFTTD) from the Census Bureau The Longitudinal Firm Trade Transactions Database (LFTTD), developed by the U.S. Census Bureau, links merchandise export and import transactions from U.S. Customs data to firms in the Census Business Register. Covering data from 1992 onward, LFTTD provides detailed shipment-level information—including product codes (at the HS-10 level), trade value and quantity, partner country, transport mode, and port of entry or exit—matched to firm-level identifiers using a combination of deterministic (EIN-based) and probabilistic (name-based and machine learning) algorithms (Kamal and Ouyang, 2020).

In addition to the key information described above, we also utilize trade cost mea-


asures available in the LFTTD since the early 2000s, including shipment-level shipping costs (freight and insurance charges) and duties paid. We aggregate these shipment-level costs to the firm-month level by computing value-weighted averages across all of a firm’s import transactions in a given month. Specifically, for firm j in month t , we define ad valorem shipping costs as the ratio of total freight and insurance charges to total customs value across all of the firm’s import shipments.

Figure 1 plots the value-weighted average shipping costs as a percentage of shipment value from 2001 to 2022, broken out by major source countries. As shown in the figure, before the pandemic, the average shipping cost hovered around 4.5 percent. During the pandemic, however, it jumped to over 8 percent. Importantly, the surge was broad-based: all major trade partner countries experienced rising shipping costs, though with some heterogeneity in magnitude. Shipping costs for imports from China and India rose particularly sharply, reflecting the severe congestion at trans-Pacific and Indian Ocean ports.

Producer Price Index (PPI) data from the BLS The Producer Price Index (PPI), published by the U.S. Bureau of Labor Statistics, measures the average change over time in the prices received by domestic producers for their goods and services. The BLS collects price data from a large sample of establishments on a monthly basis, covering a wide range of commodities and services. The confidential microdata underlying the published PPI indices contain information on individual price quotes, product characteristics, industry classifications (at the 6-digit NAICS level), and establishment identifiers. This level of granularity allows us to construct firm-level price indices that control for product composition effects.

Following Gilchrist et al. (2017), we use the micro data available from the BLS to construct firm-level inflation measures as the weighted average of price changes across

FIGURE 1: Trade Costs



NOTE: This figure plots value-weighted average ad valorem shipping costs (freight and insurance charges as a percentage of customs value) for U.S. merchandise imports from 2001 to 2022, disaggregated by major source country. The bold line denotes the average across all source countries; individual lines correspond to China, Germany, India, Italy, Japan, Korea, and the rest of the world (RoW). Source: Authors' calculations based on LFTTD data.


goods within each firm:

$$\Delta p_{f,t} = \frac{1}{n_f} \sum_{k=1}^{n_f} w_{k,f,t} \Delta p_{k,f,t} \quad (1)$$

where $\Delta p_{k,f,t}$ denotes log price changes for each good k in firm f for a specific period t , and n_f is the number of goods produced by the firm. The weights $w_{k,f,t}$ are based on the relative importance weights assigned by the BLS and on firm-level sales data.¹

¹See Gilchrist et al. (2017) for a full description of how to construct firm-level price indices from the PPI data.

FIGURE 2: 12-month PPI Inflation



NOTE: This figure plots the 12-month PPI inflation rate for U.S. manufacturing firms, as published by the BLS. The shaded band marks the COVID-19 pandemic. Source: BLS.

Figure 2 displays the 12-month PPI inflation rate for manufacturing firms, as published by the BLS. As illustrated in the figure, by June 2022, 12-month PPI inflation for manufacturers approached 20 percent, reflecting the combined effects of supply chain disruptions, rising energy and commodity prices, and strong demand.

Merged Dataset To examine how trade costs affect firms' pricing behavior, we construct a novel dataset by merging micro-level data from the LFTTD and PPI databases. The merge is based on a multi-step name-matching algorithm described in detail in

Appendix A. Because the two databases do not share a common firm identifier, we employ a combination of exact matching, word-based matching, and fuzzy matching procedures to link business names across the two data sources. We impose several quality filters to minimize false matches, including requiring consistency in industry classification and geographic location.

The resulting sample is a panel of approximately 6,000 U.S. manufacturing importing firms observed at a monthly frequency from 2005 to 2022. These firms span a broad range of manufacturing industries, including food products, chemicals, metals, machinery, electrical equipment, and transportation equipment. Summary statistics for the merged sample are presented in Table 1.

TABLE 1: Summary Statistics

	Mean	Median	Std. Dev.	P10	P90
<i>Panel A: Trade Cost Variables</i>					
Ad valorem shipping costs (%)	4.82	3.67	4.15	1.02	9.85
Number of source countries	5.3	4	4.1	1	11
Number of trade routes	7.8	6	5.9	2	16
<i>Panel B: Producer Price Variables</i>					
Monthly log price change (%)	0.12	0.00	2.84	−2.15	2.38
12-month log price change (%)	1.45	0.87	7.62	−5.21	8.73
Number of products per firm	3.7	3	2.8	1	7

NOTE: This table reports summary statistics for the merged LFTTD–PPI sample, which includes approximately 6,000 U.S. manufacturing importing firms observed at a monthly frequency from 2005 to 2022. Ad valorem shipping costs and duties are expressed as a percentage of customs value. Source: Authors’ calculations.

2.2 Empirical Specification

We now turn to estimating the pass-through of trade costs to U.S. producer prices.

To assess the impact of trade costs, we regress future producer prices (s months

ahead) on current shipping costs, including an interaction term that allows the relationship to differ during the pandemic.

Our baseline empirical specification is as follows:

$$\ln p_{j,k,t+s} = \gamma \ln c_{j,t} + \gamma_{crisis} \text{Crisis}_t \cdot \ln c_{j,t} + \sum_{l=1}^L \beta_l \ln p_{j,k,t-l} + FE + \epsilon_{j,k,t+s}, \quad (2)$$

where $\ln p_{j,k,t+s}$ denotes the log producer price of product k produced by firm j in month $t + s$, and $\ln c_{j,t}$ represents the log ad valorem shipping costs faced by firm j in month t . The indicator variable Crisis_t equals one from March 2020 through December 2022—the pandemic period—and zero otherwise; no other episode enters it. The interaction term $\text{Crisis}_t \cdot \ln c_{j,t}$ allows pass-through to differ during the pandemic.

The coefficient γ captures the average pass-through of shipping costs to producer prices during normal times. The coefficient γ_{crisis} captures the *incremental* pass-through during the pandemic, so that total pandemic-period pass-through is $\gamma + \gamma_{crisis}$. We estimate the model at multiple horizons ($s = 3, 6, 9, 12$ months) to trace out the dynamics of cost pass-through.

To control for unobservable shocks that are common within the same industry and month, we include industry-by-month fixed effects at the 3-digit NAICS level. These fixed effects absorb industry-wide cost and demand shocks, ensuring that our estimates are identified from variation in shipping costs *across firms within the same industry and month*. In our full specification, we further include firm fixed effects to identify pass-through from within-firm variation in trade costs over time. Lastly, we control for price dynamics by including up to $L = 12$ monthly lags of producer prices at the firm-product level, thereby accounting for persistence or stickiness in price setting.

Standard errors are clustered at the firm level throughout to allow for arbitrary serial correlation in the error term within firms.

2.3 Results: Shipping Cost Pass-through

We begin by examining the pass-through of shipping costs to producer prices. Table 2, Panel A reports OLS estimates of Equation (2), while Panel B presents IV estimates using plausibly exogenous variation in shipping costs. Across specifications, we estimate price responses at 3-, 6-, 9-, and 12-month horizons.

2.3.1 OLS Estimates

The baseline coefficient on shipping costs, γ , is small and statistically insignificant in the OLS specifications. In columns (1)–(4), which omit firm fixed effects, the estimated pass-through is close to zero at all horizons. When we include firm fixed effects in columns (5)–(8), the coefficients become slightly negative but remain statistically indistinguishable from zero. These results suggest that, on average, shipping costs do not significantly affect producer prices outside of crisis periods.

By contrast, the interaction term, γ_{crisis} , which captures the incremental effect of shipping costs during the COVID-19 pandemic, is consistently positive and statistically significant. However, the OLS estimates suggest economically modest effects. For example, at the 12-month horizon, a 10 percent increase in shipping costs during the pandemic is associated with only a 0.029 percent increase in producer prices when firm fixed effects are not included (column (4)), and a 0.024 percent increase when firm fixed effects are included (column (8)). The estimated effect grows monotonically with the horizon: under the firm fixed effects specification, γ_{crisis} rises from 0.0006 at three months to 0.0013, 0.0015, and 0.0024 at the six-, nine-, and twelve-month horizons, respectively.

TABLE 2: Pass-through of shipping costs

A. OLS								
horizon (month)	3	6	9	12	3	6	9	12
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
γ	0.0001 (0.0001)	0.0001 (0.0002)	0.0003 (0.0003)	0.0003 (0.0004)	-0.0001 (0.0002)	-0.0002 (0.0003)	-0.0002 (0.0004)	-0.0004 (0.0004)
γ_{crisis}	0.0010 (0.0004)	0.0016 (0.0007)	0.0020 (0.0010)	0.0029 (0.0014)	0.0006 (0.0005)	0.0013 (0.0007)	0.0015 (0.0010)	0.0024 (0.0014)
Time $\times$ Ind. F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Firm F.E.	N	N	N	N	Y	Y	Y	Y
R^2	0.92	0.87	0.83	0.78	0.92	0.88	0.84	0.81
Obs.	947,000	907,000	867,000	821,000	947,000	907,000	867,000	821,000
B. IV								
horizon (month)	3	6	9	12	3	6	9	12
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
γ	0.0040 (0.0009)	0.0075 (0.0016)	0.0104 (0.0022)	0.0136 (0.0029)	0.0066 (0.0123)	0.0130 (0.0197)	0.0170 (0.0280)	0.0216 (0.0356)
γ_{crisis}	0.0115 (0.0043)	0.0201 (0.0076)	0.0273 (0.0107)	0.0368 (0.0144)	0.0112 (0.0045)	0.0208 (0.0080)	0.0279 (0.0110)	0.0357 (0.0144)
Time $\times$ Ind. F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Firm F.E.	N	N	N	N	Y	Y	Y	Y
R^2	0.91	0.85	[h9]	0.74	0.88	0.80	0.73	0.66
Obs.	942,000	902,000	862,000	817,000	942,000	902,000	862,000	817,000

NOTE: The dependent variable in all specifications is $\ln p_{j,k,t+s}$ for a horizon s of 3, 6, 9, and 12 months. The key explanatory variable, $c_{j,t}$, is ad valorem shipping costs for firm j in month t . Crisis _{t} is a dummy variable equal to one between March 2020 and December 2022 (the pandemic period), and zero otherwise. All specifications include 3-digit NAICS industry-by-month fixed effects and up to 12 lags of the dependent variable. Standard errors, clustered at the firm level, are reported in parentheses.

2.3.2 Bartik-Style Instruments

One potential reason for the small magnitude of the OLS estimates is downward bias due to demand factors. Negative demand shocks can simultaneously reduce both import demand (thereby lowering measured shipping costs through compositional effects or reduced congestion) and producer prices, generating a spurious positive correlation between shipping costs and prices that attenuates the estimated pass-through. To better isolate exogenous supply-side variation in trade costs, we construct Bartik-style instrument variables that exploit plausibly exogenous shifts in shipping costs, while mitigating endogeneity concerns arising from demand shocks.

Specifically, our instruments are formed by interacting a firm's historical exposure to different trade routes with route-level national average shipping costs. To avoid

mechanical correlation between the instrument and the firm’s own outcomes, we exclude all firms in the same 2-digit NAICS industry as firm j when calculating national average shipping costs for firm j . Formally, for firm j in month t , its Bartik-style instrument is defined by

$$B_{j,t} = \sum_r z_{j,r,t} C_{r,t}, \quad (3)$$

where

$$\begin{aligned} z_{j,r,t} &= \frac{\sum_{t' \in [t-60, t-1]} \text{value}_{j,r,t'}}{\sum_{r'} \sum_{t' \in [t-60, t-1]} \text{value}_{j,r',t'}}, \\ C_{r,t} &= \frac{\sum_{k \notin \text{Ind}(j)} \text{cost}_{k,r,t}}{\sum_{k \notin \text{Ind}(j)} \text{value}_{k,r,t}}. \end{aligned}$$

Here, $z_{j,r,t}$ denotes the share of firm j ’s total international trade routed through trade route r during the 60 months preceding month t , capturing the firm’s predetermined exposure to route-level shipping costs. The variable $C_{r,t}$ is the value-weighted average ad valorem shipping costs on route r in month t , computed using data from all firms outside firm j ’s 2-digit NAICS industry (denoted $\text{Ind}(j)$). The 60-month lag structure for $z_{j,r,t}$ ensures that our instruments are based on pre-determined trade patterns, leveraging plausibly exogenous variation in shipping costs.

Two identifying routes are available for shift-share designs, and it is worth stating which one we rely on. [Goldsmith-Pinkham et al. \(2020\)](#) show that a shift-share instrument is numerically equivalent to using the exposure shares as instruments, so that identification rests on the exogeneity of the *shares*. [Borusyak et al. \(2022\)](#) establish a complementary route in which identification comes instead from quasi-random *shocks*, with consistency delivered by many shocks that are as good as randomly assigned conditional on exposure. Our design is closer to the second: the shifts $C_{r,t}$ are leave-one-out national route-level shipping costs constructed excluding the firm’s

own 2-digit industry, and the panel contains many route-month shocks whose variation is driven by shipping capacity, port congestion, and vessel deployment on particular lanes—plausibly unrelated to the demand conditions facing any individual U.S. manufacturer.

The two assumptions are not mutually exclusive, and our design also speaks to the shares. The identifying assumption in that reading is that, conditional on industry-by-time fixed effects and firm fixed effects, a firm’s predetermined pattern of trade route utilization is uncorrelated with future product-level demand shocks that affect its producer prices. This assumption would be violated if, for example, firms that historically relied on certain trade routes also experienced systematically different demand trends. The inclusion of industry-by-time fixed effects helps mitigate this concern by absorbing industry-level demand shocks, and the leave-one-out construction of $C_{r,t}$ ensures that the instrument is not mechanically driven by the firm’s own shipping cost realizations.

2.3.3 IV Estimates

Table 2, Panel B presents IV estimates using the Bartik-style instruments described above. By mitigating the downward bias arising from demand shocks, the IV estimates yield substantially larger coefficients relative to the OLS results. The baseline pass-through outside the crisis period (γ) is positive and statistically significant when firm fixed effects are not included (columns (1)–(4)). However, once firm effects are included (columns (5)–(8)), the coefficients become statistically insignificant. The contrast underscores the importance of controlling for unobserved firm-level heterogeneity. The significant and positive coefficients in columns (1)–(4) likely reflect between-firm relationships—firms with higher shipping costs may also have higher producer prices on average, perhaps due to differences in product quality, reliance on imported intermediate inputs, or other persistent firm characteristics. In contrast, the lack of significance

once firm fixed effects are included suggests that there is little evidence of *within-firm* pass-through of shipping costs to producer prices outside crisis periods.

Crucially, the interaction term (γ_{crisis}) remains strongly statistically significant and is an order of magnitude larger than the OLS estimates. Its magnitude increases with the horizon, rising from approximately 0.011 at the 3-month horizon to 0.021, 0.028, and 0.036 at the six-, nine-, and twelve-month horizons, respectively, regardless of whether firm fixed effects are included. These estimates indicate that the pass-through of shipping costs to producer prices was substantially amplified during the COVID-19 pandemic.

To put the magnitudes in perspective, the average ad valorem shipping cost increased from about 4.5 percent to over 8 percent during the pandemic, an approximate 75 percent increase. The estimates imply that a firm whose shipping costs rose by 10 percent more than its industry peers raised its price by about 0.36 percent more after twelve months. Because the coefficients are identified from cross-firm variation within industry-months, they cannot be converted into an aggregate inflation contribution by scaling them up by the aggregate shipping move; we compute the model-implied aggregate contribution in Section 4.

2.3.4 Discussion of Results

Overall, the results point to limited pass-through of shipping costs to producer prices in normal times, but significantly stronger effects during the pandemic-era supply-chain disruption. The IV estimates further suggest that accounting for endogeneity in shipping costs reveals economically meaningful and delayed price responses, with pass-through rising over time.

The sharp contrast between normal times and the pandemic period is consistent with a mechanism in which the effectiveness of firms' input-adjustment margins de-

depends on the breadth of the underlying shock. In normal times, shipping-cost fluctuations are more idiosyncratic across routes and source countries. Firms can therefore partly absorb these shocks by adjusting the mix of imported and domestic inputs for affected goods, limiting their effect on marginal cost and output prices. During the pandemic, by contrast, shipping costs rose broadly across many routes and source countries. This broad-based increase reduced the effectiveness of the foreign-domestic input-substitution margin, allowing a larger share of the shock to pass through into marginal costs and producer prices.

This interpretation motivates the quantitative model below. The model does not rely on firms switching across foreign suppliers or routes. Instead, it holds the set of imported varieties and routes fixed and allows firms to adjust the import intensity of intermediate inputs. Combined with menu costs, this import-intensity margin helps explain why normal-times pass-through is close to zero, while broad-based pandemic-era shipping shocks generate larger and more persistent price responses.

3 A Menu-Cost Model of Shipping-Cost Pass-Through

To understand the empirical results—negligible pass-through of shipping costs in normal times, and pass-through an order of magnitude larger during the pandemic—we develop a menu-cost model in the tradition of [Goloso and Lucas \(2007\)](#), [Nakamura and Steinsson \(2008\)](#), and [Midrigan \(2011\)](#), building most directly on the price-gap framework of [Gagliardone et al. \(2025\)](#), augmented with imported intermediate inputs whose price depends on shipping costs.

The model rationalizes state-dependent pass-through through an import-intensity margin: firms can adjust the mix of imported and domestic inputs within a fixed set of imported varieties and routes. In normal times, this foreign-domestic input-

substitution margin helps absorb idiosyncratic shipping-cost shocks before they fully reach marginal cost. During the pandemic, shipping-cost shocks were broad-based and persistent, reducing the effectiveness of this margin and generating larger price responses. We develop the model in full here and calibrate it in Section 4.

3.1 Households and demand

A representative household has constant-elasticity-of-substitution preferences over a continuum of differentiated varieties produced by firms $f \in [0, 1]$,

$$Y_t = \left[\int_0^1 Y_t(f)^{(\sigma-1)/\sigma} df \right]^{\sigma/(\sigma-1)}, \quad \sigma > 1, \quad (4)$$

which yields the familiar isoelastic demand and price index

$$Y_t(f) = \left(\frac{P_t(f)}{P_t} \right)^{-\sigma} Y_t, \quad P_t = \left[\int_0^1 P_t(f)^{1-\sigma} df \right]^{1/(1-\sigma)}. \quad (5)$$

3.2 Production and the role of shipping costs

Each firm f produces with labor $L_t(f)$ and a composite intermediate input $X_t(f)$ via Cobb–Douglas technology,

$$Y_t(f) = A_t(f) X_t(f)^\phi L_t(f)^{1-\phi}, \quad (6)$$

where $A_t(f)$ is firm-specific productivity, with $a_t(f) \equiv \log A_t(f)$ following a random walk, and $\phi \in (0, 1)$ is the intermediate input share. The intermediate composite aggregates N origin-specific inputs with Cobb–Douglas weights $\gamma_c(f)$,

$$X_t(f) = \prod_{c=1}^N X_{c,t}(f)^{\gamma_c(f)}, \quad \sum_{c=1}^N \gamma_c(f) = 1, \quad (7)$$

and within each origin c , domestic ($Z_{c,t}$) and foreign ($M_{c,t}$) varieties combine through a CES bundle with elasticity ζ ,

$$X_{c,t}(f) = [Z_{c,t}(f)^{\zeta/(1+\zeta)} + a_c^{1/(1+\zeta)} M_{c,t}(f)^{\zeta/(1+\zeta)}]^{(1+\zeta)/\zeta}, \quad (8)$$

with $a_c > 0$ an efficiency weight on foreign varieties. The elasticity ζ controls how aggressively firms substitute between imported and domestic inputs when delivered foreign input prices rise. We refer to this as the foreign-domestic input-substitution margin, or equivalently the import-intensity margin. For notational simplicity, we suppress the firm argument below and write γ_c for the Cobb–Douglas weights $\gamma_c(f)$.

The price of the domestic variety is denoted $P_{c,t}^H$. The foreign variety is subject to an iceberg shipping cost: delivering one unit costs $\tau_{c,t} \geq 1$ units, so its delivered price is $P_{c,t}^F \tau_{c,t}$, where $P_{c,t}^F$ is the foreign producer price of the same origin- c input. Shipping costs $\tau_{c,t}$ are the empirical object $c_{j,t}$ of Section 2. Cost minimization within the origin- c bundle implies that the share of expenditure on imports $\alpha_{c,t}$ —which is governed by the relative price $R_{c,t} \equiv P_{c,t}^F/P_{c,t}^H$ and the shipping cost—is given by

$$\alpha_{c,t} = \frac{a_c (R_{c,t} \tau_{c,t})^{-\zeta}}{1 + a_c (R_{c,t} \tau_{c,t})^{-\zeta}}. \quad (9)$$

Defining the origin- c cost-reduction factor $b_{c,t} = [1 + a_c (R_{c,t} \tau_{c,t})^{-\zeta}]^{1/\zeta}$ and the aggregate factor $B_t(f) = \prod_c b_{c,t}^{\gamma_c(f)}$, the firm's nominal marginal cost is

$$MC_t^m(f) = \frac{C_t^*(f)}{A_t(f) B_t(f)^\phi}, \quad (10)$$

where $C_t^*(f) = (P_t^H(f)/\phi)^\phi (W_t/(1-\phi))^{1-\phi}$ is the unit cost of a hypothetical non-importer and $P_t^H(f) = \prod_c (P_{c,t}^H/\gamma_c(f))^{\gamma_c(f)}$.

Equation (10) makes the shipping channel explicit. By Shephard's lemma, the

elasticity of the origin- c cost index to the delivered input price equals the import expenditure share, $\partial \log b_{c,t} / \partial \log \tau_{c,t} = -\alpha_{c,t}$, so a rise in $\tau_{c,t}$ raises marginal cost with elasticity

$$\frac{\partial mc_t^n(f)}{\partial \log \tau_{c,t}} = \phi \gamma_c(f) \alpha_{c,t}, \quad (11)$$

where $mc_t^n(f) \equiv \log MC_t^n(f) = \log C_t^*(f) - a_t(f) - \phi \log B_t(f)$ is the logarithm of the nominal marginal cost. The shipping-cost to marginal-cost elasticity is thus a product of the intermediate-input share ϕ , the origin weight γ_c , and the import share $\alpha_{c,t}$. The pass-through into cost is thus small precisely when the import share is small, and—because the firm substitutes away from the dearer origin (9)—the share is itself decreasing in $\tau_{c,t}$, buffering the cost increase.

Summing over origins, the elasticity to a *common* shock is $\phi s_t(f)$ with $s_t(f) \equiv \sum_c \gamma_c(f) \alpha_{c,t}$ the aggregate import share. Let $\bar{s}(f) \equiv \sum_c \gamma_c(f) \bar{\alpha}_c$ denote its steady-state value.

3.3 Shipping costs and the static target price

Absent any adjustment friction, the firm would set the price that maximizes flow profit, the *static target price* $p_t^o(f)$. Following Nakamura and Steinsson (2008), we assume that the nominal wage satisfies $W_t = \frac{\sigma-1}{\sigma} P_t$ and the real unit cost of a non-importer, scaled by the steady-state value of the cost-reduction factor, is constant and equal to $\frac{C_t^*}{P_t} = \frac{\sigma-1}{\sigma} \bar{B}^\phi$. Under this assumption together with CES demand and the marginal cost (10), the static target price is a constant markup over marginal cost, which in logs becomes

$$p_t^o(f) = \log \frac{\sigma}{\sigma-1} + mc_t^n(f) = \log P_t - z_t(f), \quad (12)$$

where $z_t(f)$ is the firm's combined cost shifter, given by

$$z_t(f) \equiv a_t(f) + \phi \log(B_t(f)/\bar{B}(f)). \quad (13)$$

The target moves one-for-one with the aggregate price $\log P_t$ and falls with the firm's productivity and import-cost advantage $z_t(f)$. Innovations to $z_t(f)$ combine productivity shocks, relative import-price shocks, and shipping-cost shocks. Specifically, using hats to denote log-deviations from steady state, $\hat{B}_t = \log B_t - \log \bar{B}$, we show:

$$\hat{B}_t(f) = - \sum_{c=1}^N \gamma_c(f) \bar{\alpha}_c \left[\hat{R}_{c,t} + \hat{\tau}_{c,t} \right]. \quad (14)$$

We take the relative price $R_{c,t} = \bar{R}_c$ as fixed (i.e., $\hat{R}_{c,t} = 0$), so the only variation is in shipping costs. Let $\Delta\alpha_{c,t} \equiv \alpha_{c,t} - \bar{\alpha}_c$ denote the first-order level deviation of the share from steady state, then we can show $\frac{\partial \log \alpha_{c,t}}{\partial \log \tau_{c,t}} = -\zeta(1 - \alpha_{c,t})$ and

$$\Delta\alpha_{c,t} = \bar{\alpha}_c \cdot (-\zeta(1 - \bar{\alpha}_c)) \hat{\tau}_{c,t} = -\zeta \bar{\alpha}_c (1 - \bar{\alpha}_c) \hat{\tau}_{c,t}, \quad (15)$$

implying that the firm reallocates expenditure away from the origin whose shipping cost has risen, at a rate increasing in ζ .

We now model the shipping costs. Up until $t = 0$, the economy has been in steady state and the crisis occurs starting $t = 1$ and ending $t = T_R$. We assume that during the crisis period, each route-level shipping deviation can be decomposed into an origin-specific and a common component,

$$\hat{\tau}_{c,t} = v_{c,t} + u_t, \quad (16)$$

where $v_{c,t}$ is origin-specific (common to all firms), and $u_t \equiv \log(\tau_t/\bar{\tau})$ is common across

origins.

In normal times, we assume $u_t = 0$ (i.e., $\hat{\tau}_{c,t} = v_{c,t}$) and each firm has diversified origins such that $\sum_c \gamma_c(f) \bar{\alpha}_c v_{c,t} \approx 0$. As a result, the firm-level cost shifter decomposes as

$$z_t(f) = \tilde{z}_t(f) - \phi s(f) u_t, \quad (17)$$

where $\tilde{z}_t(f) \equiv a_t(f) - \phi \sum_c \gamma_c(f) \bar{\alpha}_c [\hat{R}_{c,t} + v_{c,t}] = a_t(f)$ and $s(f) \equiv \sum_{c=1}^N \gamma_c(f) \bar{\alpha}_c$ is the firm's steady-state aggregate import share. Following GGLT, we assume that both the aggregate price and the idiosyncratic component of the cost shifter follow random walk:

$$\log P_t = \mu + \log P_{t-1} + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2) \text{ i.i.d.} \quad (18)$$

and

$$\tilde{z}_t(f) = \tilde{z}_{t-1}(f) + \tilde{\epsilon}_t(f), \quad \tilde{\epsilon}_t(f) \sim N(0, \sigma_\epsilon^2),$$

with $\tilde{\epsilon}_t(f)$ i.i.d. across firms and time, combining productivity innovations and idiosyncratic import-cost innovations.

Importantly, the crisis path of the common component $\{u_t\}_{t=0}^{T_R}$ is modeled as a MIT shock. That is, the crisis path $\{u_t\}_{t=0}^{T_R}$ is deterministic, is assigned probability zero by firms for $t < 0$, and is common knowledge from $t = 0$ onward (perfect foresight). Firms' beliefs about $\tilde{\epsilon}$ and η are unchanged. Policy functions are time-dependent during the window $t \in [0, T_R]$ and are computed by backward induction from the stationary solution, to which the economy reverts for $t > T_R$.

3.4 Menu cost and the firm's dynamic problem

Let $P_t(f)$ denote the price the firm charges at t . Real profits before adjustment costs depend on the firm's state only through the relative price $P_t(f)/P_t$ and the cost shifter

$z_t(f)$: using the demand function and $MC_t^n(f) = \frac{\sigma-1}{\sigma} P_t e^{-z_t(f)}$ from (12),

$$\begin{aligned} \Pi_t \left(\frac{P_t(f)}{P_t}, z_t(f) \right) &\equiv \left(\frac{P_t(f)}{P_t} \right)^{-\sigma} Y \frac{P_t(f) - MC_t^n(f)}{P_t} \\ &= Y e^{(\sigma-1)z_t(f)} \left[e^{(1-\sigma)\tilde{p}_t(f)} - e^{-\bar{\mu}-\sigma\tilde{p}_t(f)} \right], \end{aligned} \quad (19)$$

where $\tilde{p}_t(f) \equiv p_t(f) - p_t^o(f)$ is the log deviation of the price from the static target and $\bar{\mu} \equiv \log \frac{\sigma}{\sigma-1}$ is the log markup. We show in the appendix (see Lemma 1) that

$$\Pi_t \left(\frac{P_t(f)}{P_t}, z_t(f) \right) \approx \bar{\Pi}_t(z_t(f)) \left(1 - \frac{\sigma(\sigma-1)}{2} \tilde{p}_t(f)^2 \right), \quad (20)$$

where $\bar{\Pi}_t(z_t(f)) \equiv \frac{Y}{\sigma} e^{(\sigma-1)z_t(f)}$ is the gap-zero profit level, the profit at $\tilde{p}_t(f) = 0$; it depends on the firm only through its cost shifter.

Price adjustment is subject to a stochastic menu cost, which is proportional to the firm's profit scale: $\chi_t(f) = \tilde{\chi}_t(f) \bar{\Pi}_t(z_t(f))$, with $\tilde{\chi}_t(f) \sim U(0, \bar{\chi})$ i.i.d. across firms and time.

The timing within a period is as follows: the firm observes $z_t(f)$ and the aggregate state, draws the menu cost $\chi_t(f)$, and then chooses whether to adjust, $I_t(f) \in \{0, 1\}$, and, if it adjusts, the new price $P_t(f)$; if it does not adjust, $P_t(f) = P_{t-1}(f)$.

Let $x_{t-1}(f) \equiv p_t^o(f) - p_{t-1}(f)$ be the firm's price gap (the distance between the desired and the inherited price). From the target price (12), $x_{t-1}(f) = -\log(P_{t-1}(f)/P_t) - z_t(f)$, a function of the two state variables. If the firm does *not* adjust at t ,

$$x_t(f) = x_{t-1}(f) + w_{t+1}(f), \quad (21)$$

where $w_{t+1} \equiv p_t^o(f) - p_{t-1}^o(f) = \mu + \eta_t - \tilde{\epsilon}_t(f) + \phi s(f) \Delta u_t$ denotes the increment in the static target price. If it adjusts to the reset price $p_t^*(f)$, we denote the post-decision gap by $x_t^* \equiv p_t^o(f) - p_t^*(f)$, which we refer to as the *optimal reset gap*, and the gap

entering $t + 1$ is $x_t^* + w_{t+1}(f)$. In normal times w is i.i.d.; during the window it carries the deterministic, perfectly foreseen component $\phi s(f) \Delta u$.

The firm's problem is summarized by three value functions: $V_t(x_{t-1}(f))$, the value of a firm that enters the period with gap $x_{t-1}(f)$ and does *not* adjust; V_t^a , the value attained by an adjusting firm at its optimal reset gap, before netting out the menu cost it pays; and $\mathcal{E}_t(x_{t-1}(f))$, the expected value at the start of the period, before the menu-cost draw, under the optimal adjustment decision. Their level counterparts, $\mathcal{V}_t^N$, $\mathcal{V}_t^A$, and $\mathcal{V}_t$, are stated in Appendix C as equations (39)–(41). In parallel to those, the scaled value functions satisfy

$$V_t(x_{t-1}(f)) = -\frac{\sigma(\sigma-1)}{2}x_{t-1}(f)^2 + \beta E_t[\mathcal{E}_{t+1}(x_t(f))], \quad (22)$$

$$V_t^a = \max_x \left\{ -\frac{\sigma(\sigma-1)}{2}x^2 + \beta E_t[\mathcal{E}_{t+1}(x + w_{t+1}(f))] \right\} = \max_x V_t(x), \quad (23)$$

$$\mathcal{E}_t(x_{t-1}(f)) = E_{\tilde{\chi}_t(f)} \left[\max_{I_t(f)} \{ I_t(f) [V_t^a - \tilde{\chi}_t(f)] + (1 - I_t(f)) V_t(x_{t-1}(f)) \} \right]. \quad (24)$$

Equations (22)–(24) are (39)–(41) (see the appendix) divided by $\bar{\Pi}_t(z_t(f))$ and shifted by $1/(1-\beta)$: the flow term—profits in units of the gap-zero level, net of the constant absorbed by the shift—equals $-\frac{\sigma(\sigma-1)}{2}\tilde{p}_t(f)^2$ by (20), with $\tilde{p}_t(f) = -x_{t-1}(f)$ for a firm that does not adjust, and the menu cost scales to $\tilde{\chi}_t(f)$. In (23) the adjusting firm chooses its post-decision gap x (equivalently its reset price) and faces the gap $x + w_{t+1}(f)$ next period; the second equality says that a firm that resets to gap x is thereafter in the position of a non-adjuster whose inherited gap is x , so the value of adjusting is the maximum of V_t over the gap, attained at the optimal reset gap x_t^* , and an adjusting firm obtains V_t^a less the menu cost it pays. Lemma 2 in the appendix gives (24) in closed form, $\mathcal{E}_t(x) = V_t(x) + G_t(x)^2/(2\bar{\chi})$ in the interior, with $G_t(x) \equiv V_t^a - V_t(x)$ the gain from adjusting; the second term is the expected value of the adjustment option net of the expected menu cost paid.

A firm adjusts when $V_t^a - \tilde{\chi}(f) \geq V_t(x_{t-1}(f))$, i.e. when the menu-cost draw falls short of the gain from adjusting $G_t(x_{t-1}(f)) = V_t^a - V_t(x_{t-1}(f))$. Because $\tilde{\chi}$ is uniform, the probability of adjustment—the *generalized hazard function* (GHF)—takes a simple closed form (see Lemma 2 in the appendix),

$$h_t(x_{t-1}(f)) = \min \left\{ \frac{V_t^a - V_t(x_{t-1}(f))}{\bar{\chi}}, 1 \right\} \approx -\frac{V''(0)}{2\bar{\chi}} x_{t-1}(f)^2 + o(x_{t-1}(f)^2). \quad (25)$$

The hazard is increasing in the absolute gap: firms with larger pricing errors are more likely to reset. This is the *selection effect*, which we will see shapes the dynamics of pass-through.

In the appendix we show that the optimal reset price p_t^* can be written as the survival-discount weighted average of expected future targets (see Lemma 3). Consequently, the gap between the optimal reset price and the static target price, denoted $\Psi_t(f)$, satisfies:

$$\Psi_t(f) \equiv p_t^*(f) - p_t^o(f) = \sum_{s \geq 0} \omega_{t,s} E_t [p_{t+s}^o(f) - p_t^o(f)],$$

where the weights $\omega_{t,s}$ are derived in the proof of Lemma 3.

It is worth noting that in GGLT, the static target price follows a random walk, so that the dynamic optimal reset price coincides with the static target: $p_t^*(f) = p_t^o(f)$. The same holds in our model before the pandemic. During the pandemic window, however, the arrival of the MIT shock $\{u_t\}_{t=0}^{T_R}$ makes the reset gap nonzero: $x_t^* = -\Psi_t(f)$, where

$$\Psi_t(f) = \mu \sum_{s \geq 0} \omega_{t,s} s + \phi s(f) \sum_{s \geq 0} \omega_{t,s} (u_{t+s} - u_t), \quad (26)$$

is a *deterministic* sequence, common to all firms with the same exposure.

3.5 Model Implications

Normal times: why pass-through is negligible. In normal times ($u_t = 0$), the route-level deviations $v_{c,t}$ are idiosyncratic across origins and transitory, and—as assumed above—firms’ route portfolios are sufficiently diversified that the portfolio average $\sum_c \gamma_c(f) \bar{\alpha}_c v_{c,t}$ is negligible: idiosyncratic shipping movements wash out of the firm-level cost shifter, and the substitution response (15) shrinks what survives further—though only modestly, since the share responds to a cost movement only at second order in its size (Section 4). The model’s account of the insignificant normal-times estimates in Table 2 is therefore deliberately modest: prices do not respond to normal-times shipping variation because, after diversification and substitution, essentially none of it reaches marginal cost. The within-firm variation that remains in the *measured* rate—finite portfolios, shipment composition, and invoicing timing—is treated as orthogonal to marginal cost.

The crisis: a common, persistent, foreseen shock. The defining feature of the pandemic was that shipping costs rose *simultaneously across virtually all routes*—a large, common shock rather than the idiosyncratic noise of normal times. Diversification then offers no protection, and the substitution margin, while still operating within each origin bundle, offsets the common increase only partially: by (17), the common component enters the firm’s cost shifter with loading $\phi s(f)$. What remains is the question of how much of the shock *prices* absorb, and when. The next result answers it in closed form.

The firm’s import-value weight on route c is its steady-state share of route c in imported-input expenditure,

$$w_c(f) \equiv \frac{\gamma_c(f) \bar{\alpha}_c}{\sum_{c'} \gamma_{c'}(f) \bar{\alpha}_{c'}} = \frac{\gamma_c(f) \bar{\alpha}_c}{s(f)}, \quad \sum_c w_c(f) = 1, \quad (27)$$

predetermined at steady-state values—the model counterpart of the shipment-value weights in the empirical charges-per-value construction. The firm’s measured ad valorem rate is the value-weighted average of its route rates, $\tau_t(f) = \sum_c w_c(f) \tau_{c,t}$.

Let $\bar{\tau}(f) \equiv \sum_c w_c(f) \bar{\tau}_c$ and define the multiplier-value weights $\omega_c(f) \equiv w_c(f) \bar{\tau}_c / \bar{\tau}(f)$, $\sum_c \omega_c(f) = 1$. Then we obtain the following log-linearization around steady state,

$$\hat{\tau}_t(f) \equiv \ln \frac{\tau_t(f)}{\bar{\tau}(f)} = \ln \sum_c \omega_c(f) e^{u_t + v_{c,t}} \approx \sum_c \omega_c(f) (u_t + v_{c,t}) \approx u_t. \quad (28)$$

The small idiosyncratic component is orthogonal to the deterministic crisis path and plays no role in the derivative below, so we drop it and keep the common component.

For t and $t + h$ inside the crisis window, the *level pass-through coefficient* is the first-order effect of the current multiplier on the expected future price

$$\gamma_{crisis}(h; t) \equiv \frac{dE_t[p_{t+h}(f)]}{d\hat{\tau}_t(f)}, \quad (29)$$

where the expectation averages over the firm’s idiosyncratic shocks and adjustment history. Because the crisis path is deterministic, the derivative is taken *along the path*—scaling the firm’s entire path proportionally—which is precisely the variation that cross-firm differences in shipping exposure identify in the data. Equation (29) is a causal object: no regression enters its definition.

Proposition 1. *Under the constant-hazard approximation with monthly survival probability θ (as defined in Section 3.5), for any MIT path $\{u_t\}$ and any t, h with $t \geq 1$,*

$$\gamma_{crisis}(h; t) = \phi_s(f) \frac{\theta^h \ell_t + \sum_{j=1}^h (1 - \theta) \theta^{h-j} \varphi(t + j)}{u_t} = \phi_s(f) \frac{\ell_{t+h}}{u_t}, \quad (30)$$

where $\ell_\tau \equiv (1 - \theta) \sum_{j \geq 0} \theta^j \varphi(\tau - j) \mathbf{1}\{\tau - j \geq 1\}$, $\varphi(a) \equiv \sum_{s \geq 0} \omega_s u_{a+s}$, and $\omega_s = (1 - \beta\theta)(\beta\theta)^s$.

Proposition 1 has a transparent reading. The loading $\varphi(a)$ is what a price reset at date a builds in—the survival-discount forward average of the remaining path—and the *incorporation level* ℓ_τ is the probability-weighted average of these loadings over past reset dates: the fraction of the shock embedded in the average price at τ . Crisis pass-through at horizon h is simply the incorporation level at $t + h$, relative to the current shock. Two implications follow immediately. The horizon profile of pass-through is the incorporation path itself, read forward from the projection date—so its *shape* identifies the shape of the underlying cost path. And because φ is forward-looking, anticipated future increases are built into prices *before* they occur, while movements expected to die out quickly are heavily discounted. In the next section we evaluate the formula at the calibrated parameters for a benchmark shock path and compare it with the IV estimates.

4 Calibration and Quantitative Results

We now bring the model to the data. The calibration strategy is disciplined and parsimonious: most parameters are set externally or recovered from the *first stage* of the same IV used in Section 2, leaving only two adjustment parameters to match the cross-sectional moments of price changes. We then evaluate the closed-form crisis pass-through of Proposition 1 at the calibrated parameters. A persistent shock delivers the sign and the rising horizon profile of the empirical crisis estimates; the implied magnitudes fall short of the IV estimates by a factor of three to five, a gap we trace to measured freight understating the effective delivered-cost disruption.

4.1 Calibration

The model is monthly. We set the discount factor to $\beta = 0.96^{1/12}$, the demand elasticity to $\sigma = 5.75$ and the intermediate share to $\phi = 0.65$ following standard values in the production-network literature, and the foreign cost share to $\bar{s} = 0.30$, the average imported-input share in the LFTTD. The aggregate price follows a random walk with drift, $\log P_t = \mu + \log P_{t-1} + \eta_t$ with $\eta_t \sim N(0, \sigma_\eta^2)$; we estimate $\mu = 0.00165$ and $\sigma_\eta = 0.00962$ by maximum likelihood from the manufacturing PPI over 2005–2019.

Recovering ζ from the first-stage IV. The home–foreign elasticity ζ is the single most important structural parameter. As shown above, it governs both the normal-times attenuation and the size of the common-shock pass-through. Rather than impose it, we recover it from the first stage of our Bartik instrument, which measures precisely the object ζ controls: how strongly realized input costs respond to an instrumented shipping shock once firms substitute.

Consider the identifying experiment behind the first stage: a common instrumented shipping shock of size u , holding relative prices at their steady-state values. The import share falls by $\zeta(1 - \bar{s})u$ in proportional terms, so the realized cost regressor rises *less* than the instrument by exactly that substitution wedge. The first-stage coefficient is therefore

$$\hat{\pi} = 1 - \zeta(1 - \bar{s})u. \quad (31)$$

We can read ζ directly off the estimated first stage. The own first-stage coefficient in our regression is $\Pi_{11} = 0.1719$ (the loading of the realized cost regressor on the Bartik instrument). Setting $\hat{\pi} = \Pi_{11}$, $\bar{s} = 0.30$, and $u = 0.56$ in (31) and inverting yields


$$\zeta = \frac{1 - \Pi_{11}}{(1 - \bar{s})u} = \frac{0.8281}{0.70 \times 0.56} = 2.11. \quad (32)$$

This value lies in the standard range of the trade literature and is close to the [Caliendo and Parro \(2015\)](#) estimates. Two implications are worth drawing out. First, $\zeta = 2.11$ is large enough that, by the substitution response [\(15\)](#) and the diversification argument of [Section 3.5](#), idiosyncratic shipping shocks are heavily attenuated before reaching marginal cost—rationalizing why the estimated normal-times pass-through γ is insignificant. Second, the *gap* between the IV and OLS pass-through estimates in [Table 2](#) has a structural interpretation: it reflects the same substitution wedge $\zeta(1 - \bar{s})$, since OLS conflates the exogenous cost shock with the endogenous re-sourcing response.

Calibrating the menu cost to price-change moments. Two parameters remain: the menu-cost bound $\bar{\chi}$ and the idiosyncratic shock volatility σ_ϵ . We calibrate them to match the steady-state moments of price adjustment in the manufacturing PPI micro data. [Figure 3](#) plots the monthly frequency of overall, upward, and downward price changes in [Panel \(a\)](#) and the average size of upward or downward price changes in [Panel \(b\)](#). As shown in [Panel \(a\)](#), the share of observations with any price change rises gradually from about 15–17 percent in the mid-2000s to around 20–21 percent by the late 2010s, then jumps sharply after 2020 to roughly 26–27 percent by 2022. The increase is especially pronounced for positive price changes: the frequency of upward price changes rises from roughly 10–12 percent before 2020 to about 18 percent in 2021–2022. At the same time, positive price changes shown in [Panel \(b\)](#) are consistently larger in magnitude than negative price changes. The average positive price change is usually around 6–8 percent, rising somewhat toward the end of the sample to nearly 9 percent. [Blanco et al. \(2024a\)](#) report similar findings based on CPI inflation data.

We use the pre-pandemic data between 2005 and 2019 to construct two moments, the monthly frequency of price change (0.169) and the mean absolute size of price changes (0.074), to calibrate the model parameters. Intuitively, the frequency disci-

FIGURE 3: PPI Price-Adjustment Margins



(a) Frequency of price changes

(b) Size of price changes

NOTE: This figure plots monthly frequency of overall, upward, and downward price changes in panel (a), and mean size of upward price changes in panel (b) based on BLS PPI database.

plines $\bar{\chi}$ (a higher menu cost means rarer adjustment) and the mean absolute change disciplines σ_ϵ (larger shocks open larger gaps and bigger resets). This leaves the fraction of price increases (0.594) as an untargeted moment. Solving the two-moment problem by simulated method of moments gives

$$\bar{\chi} = 0.41865, \quad \sigma_\epsilon = 0.03157. \quad (33)$$

Table 3 collects the calibration. The model matches the two targeted moments essentially exactly (frequency 0.168 versus 0.169; mean absolute change 0.074) and underpredicts the untargeted fraction of increases modestly (0.54 versus 0.59).

4.2 A jump–decay benchmark

As a benchmark, we consider the following MIT shock: a jump on impact followed by geometric decay,

$$\bar{u}_t = \rho^{t-1}, \quad t \geq 1, \quad \rho \in (0, 1), \quad (34)$$

TABLE 3: Baseline calibration

Parameter	Description	Value
<i>Externally set</i>		
β	Discount factor (monthly)	$0.96^{1/12}$
σ	Demand elasticity	5.75
ϕ	Intermediate input share	0.65
$\bar{s}$	Foreign cost share (LFTTD)	0.30
μ	Inflation drift (PPI)	0.00165
σ_η	Aggregate innovation SD (PPI)	0.00962
<i>Recovered from first-stage IV</i>		
ζ	Home–foreign trade elasticity	2.11
<i>Calibrated to price-change moments</i>		
$\bar{\chi}$	Menu-cost upper bound	0.41865
σ_ϵ	Idiosyncratic productivity SD	0.03157

NOTE: Targeted moments: monthly frequency of price change 0.169 and mean absolute price change 0.074. Model fit: 0.168 and 0.074, respectively. Untargeted fraction of price increases: 0.59 in data, 0.54 in model.

with perfect foresight from $t = 1$ onward—the impulse response of an AR(1), and a natural stand-in for a spot-freight shock that mean-reverts from the start. (The normalization of the path is immaterial: the coefficient (30) is invariant to the scale of the path, so only the shape parameter ρ matters.) This case admits a complete closed form.

Corollary 1. *For the path $\bar{u}_t = \rho^{t-1}$, the crisis pass-through coefficient is*

$$\gamma_{crisis}(h; t) = \phi s(f) \lambda(\rho) (1 - \theta) \frac{\rho^{t+h} - \theta^{t+h}}{(\rho - \theta) \rho^{t-1}}, \quad (35)$$

where

$$\varphi(a) = \lambda(\rho) \bar{u}_a, \quad \ell_\tau = \lambda(\rho) (1 - \theta) \frac{\rho^\tau - \theta^\tau}{\rho - \theta}, \quad \lambda(\rho) = \frac{1 - \beta\theta}{1 - \beta\theta\rho}.$$

The incorporation level ℓ_τ is single-peaked at $\tau^* \equiv \ln(\ln \theta / \ln \rho) / \ln(\rho / \theta)$, and the pass-

through coefficient is strictly increasing in t at every horizon, with large- t (interior) limit

$$\gamma_{crisis}(h; \infty) = \phi s(f) \lambda(\rho) (1 - \theta) \frac{\rho^{h+1}}{\rho - \theta}. \quad (36)$$

Corollary 1 is stated per unit of the log delivered-cost multiplier $\hat{\tau}_t(f)$. The empirical regressor is the log measured *rate* $\ln c_{f,t}$, so comparability with Table 2 requires multiplying every entry by the conversion factor $\varkappa(\bar{c}) = 0.043$ introduced in Section 3.5. Table 4 reports the resulting values at the calibrated parameters:

TABLE 4: Crisis pass-through under a jump-decay shock

		$h = 3$	$h = 6$	$h = 9$	$h = 12$
$\rho = 0.95$	$t = 1$ (impact month)	0.0032	0.0041	0.0042	0.0041
	$t = 6$	0.0055	0.0054	0.0050	0.0045
	$t = 12$	0.0068	0.0061	0.0054	0.0047
	$t \rightarrow \infty$ (interior)	0.0078	0.0067	0.0058	0.0049
$\rho = 0.99$	$t = 1$ (impact month)	0.0041	0.0056	0.0064	0.0067
	$t = 6$	0.0065	0.0070	0.0071	0.0072
	$t = 12$	0.0076	0.0076	0.0075	0.0074
	$t \rightarrow \infty$ (interior)	0.0082	0.0079	0.0077	0.0075

NOTE: Entries report $dE_t[p_{t+h}(f)]/d\ln c_{f,t}$ —the causal effect of a one-log-point increase in the firm’s measured ad valorem shipping rate on its expected log price h months later—under the jump-decay path with monthly persistence ρ , evaluated at the calibrated parameter values of Table 3. Rows indexed by t give the crisis month at which the projection is made ($t = 1$ is the month of impact); “interior” is the limit (36). Each entry equals equation (35) multiplied by the rate-to-multiplier conversion factor $\varkappa(\bar{c}) = 0.043$.

Because $\bar{u}_t$ is fixed at the projection date, the horizon profile of (35) is proportional to ℓ_{t+h} : *the profile in h is the incorporation path itself, read forward from t* . Hence the exact shape condition

$$\gamma_{crisis}(h; t) \text{ is increasing in } h \quad \Longleftrightarrow \quad t + h \leq \tau^*(\rho),$$

with τ^* from Corollary 1: 9.6 months at $\rho = 0.95$ and 16.6 months at $\rho = 0.99$. Two findings follow.

First, persistence delivers the empirical sign and shape. For a sufficiently persistent shock and projection dates early in the crisis, the coefficient rises monotonically over the horizons we estimate: at $\rho = 0.99$ the profile climbs from 0.0041 to 0.0067 between $h = 3$ and $h = 12$ at the impact month, and from 0.0065 to 0.0072 at $t = 6$ —qualitatively the upward-sloping pattern of the empirical local projections. The mechanism is transparent: when the shock outlives the price spell ($\rho > \theta$), incorporation is still building at the horizons we measure, so longer horizons capture more accumulated catch-up. At $\rho = 0.95$ the peak $\tau^* = 9.6$ falls inside the horizon range and the profile turns down by $h = 9$ –12; persistence near the unit root is what keeps it rising throughout.

Second, the magnitudes do not match. The empirical crisis estimates run from 0.011 to 0.036 at $h = 3$ to 12. The most favorable rising case above ($\rho = 0.99, t = 1$) falls short by a factor of 2.7 at $h = 3$, widening to 5.4 at $h = 12$. The gap widens because the model profile is not only too low but too *flat*: it rises by a factor of 1.6 from $h = 3$ to $h = 12$, against 3.3 in the data. Persistence buys the shape, in other words, but not the size or the steepness—and the shortfall cannot be closed by raising ρ . By (36), every entry is bounded by its interior limit, which approaches the static delivered-cost loading $\phi \bar{s} \varkappa(\bar{c}) = 0.0084$ as $\rho \rightarrow 1$: the causal effect of measured freight on prices can never exceed its full cost loading, and the 12-month IV estimate of 0.036 exceeds this ceiling by more than a factor of four.

4.3 The aggregate price impact

In this subsection, we compute the model’s implication on the aggregate inflation contribution. Aggregating the firm’s pricing rule over the cross-section and using the price index, inflation is the frequency-weighted average of firm-level price changes. Following GGLT, and adding the term generated by the foreseen crisis path, this

decomposes into three channels,

$$\pi_t = \underbrace{E_f[h] E_f[x]}_{\text{average gap}} + \underbrace{\text{Cov}_f(h, x)}_{\text{selection}} + \underbrace{E_f[h] \Psi_t}_{\text{anticipation}} + \text{h.o.t.}, \quad (37)$$

where h is the generalized hazard (25), x the price gap, and Ψ_t the forward-looking wedge (26). The first two terms are the decomposition in GGLT; the third is zero when the cost process is a random walk and first-order here, because firms that reset during the window price the anticipated remainder of the path.

Two properties of (37) deliver the normal-times/crisis contrast without further assumptions. In normal times the route-level shocks are idiosyncratic and mean-zero across firms, so the gap distribution does not shift, $E_f[x] = 0$, and—by symmetry— $\text{Cov}_f(h, x) \approx 0$: the aggregate impact vanishes at first order regardless of the adjustment technology. In the crisis the common component shifts every gap by the same displacement δ —the part of the delivered-cost increase not yet embedded in the average price—and all three terms activate. Evaluating them at the calibrated hazard $h(x) = h_0 + \kappa x^2$ and a symmetric gap distribution gives

$$E_f[h] = \bar{\Lambda} + \kappa \delta^2, \quad \text{Cov}_f(h, x) = 2\kappa E[x^2] \delta = 2(\bar{\Lambda} - h_0) \delta, \quad (38)$$

so the impact response of the price level to a displacement is $(3\bar{\Lambda} - 2h_0) \delta + \kappa \delta^3$, a factor of 2.60 larger than the Calvo benchmark $\bar{\Lambda} \delta$ at our calibration. The selection term is disciplined entirely by the two observables $\bar{\Lambda}$ and h_0 ; it requires no assumption about the shape of the gap distribution.

The delivered-cost increase that generates the displacement follows from Shephard's lemma applied along the path, $m^* = \int_0^{\Delta \ln \tau} \phi s(v) dv$, integrating over the size of the shock because the import share falls as shipping costs rise. With the ad valorem rate

rising from 4.5 to 8.3 percent—a 3.6 log-point increase in the delivered-cost multiplier—the import share falls from 0.300 to 0.284 at $\zeta = 2.11$, and $m^* = 0.68$ percent. The substitution margin is therefore quantitatively modest even at pandemic scale: holding the import share fixed at $\bar{s}$ would give 0.70 percent instead, so re-sourcing spares the firm about 0.02 of the 0.7 percentage points that the surge adds to marginal cost—roughly 3 percent of the increase. The share responds only at second order in the size of the cost movement, so the same is true a fortiori in normal times: what keeps normal-times pass-through near zero is the diversification of route-level shocks, not the sourcing response to them. Iterating (37) forward along the observed path of shipping costs yields the contribution to twelve-month producer-price inflation, in percentage points:

TABLE 5: Model-implied decomposition of aggregate inflation

	$E[h] E[x]$	$\text{Cov}(h, x)$	anticipation	total
Calvo benchmark	0.082	0	0.248	0.330
State-dependent adjustment	0.043	0.069	0.248	0.360

NOTE: This table reports the peak contribution of the shipping-cost surge to twelve-month PPI inflation, in percentage points, decomposed by the terms of equation (37) and evaluated along the observed path of shipping costs at the calibrated parameters of Table 3.

Two features stand out. First, the anticipation term accounts for roughly three-quarters of the total: because the surge built over nearly two years and was foreseen once underway, firms priced in increases before they occurred. Second, selection—absent by construction under Calvo—accounts for about a fifth of the total under state-dependent adjustment. The first two columns draw on the same residual gap and are not independent channels: selection speeds incorporation, which shrinks the gap that the average-gap term multiplies; their sum rises from 0.082 to 0.113 percentage points.

We read the results so far as evidence that freight and insurance charges understate the effective delivered-cost disruption of the pandemic. Port congestion, delivery delays,

expediting, and inventory costs rose together with measured freight rates; under this reading, the estimated coefficients capture the price response to the broader disruption, for which measured freight is the observable proxy, and the implied effective cost shock is several times the measured one.

5 Conclusion

This paper provides new firm-level evidence on how shipping costs pass through to U.S. producer prices, based on a novel panel that links confidential shipment-level import transactions from the LFTTD to confidential producer prices from the PPI microdata for more than 6,000 manufacturing importers over 2005–2022. The pass-through is highly state-dependent. In normal times it is statistically and economically negligible: within-firm movements in shipping costs leave producer prices essentially unchanged at every horizon. During the pandemic, when shipping costs surged broadly across routes and source countries, our Bartik-IV estimates put the incremental pass-through at 0.011 after three months, rising steadily to 0.036 after twelve—an order of magnitude above normal times, with a delayed, rising profile. The IV estimates exceed their OLS counterparts, consistent with demand shocks biasing unadjusted estimates toward zero.

A menu-cost model with imported intermediate inputs rationalizes this asymmetry and disciplines its interpretation. Idiosyncratic shipping movements largely cancel within firms’ diversified route portfolios before reaching marginal cost, while the pandemic’s common, persistent, and—once underway—foreseen surge was built into prices cumulatively through infrequent adjustment. The model reproduces the sign and rising shape of the crisis estimates in closed form, and its calibrated magnitudes imply that freight and insurance charges understate the effective delivered-cost disruption of the pandemic, which also comprised congestion, delay, and expediting costs.

The empirical results carry a clear message for assessing the inflationary consequences of trade shocks: the nature of the shock matters as much as its size. Idiosyncratic or narrowly targeted cost increases are largely absorbed by firms through sourcing adjustments and pricing frictions, whereas broad-based increases that hit many routes simultaneously, and persist, pass through to producer prices. For policymakers—and for central banks gauging the inflationary potential of supply-chain disruptions, geopolitical events, or sweeping tariff policies—this puts a premium on the breadth, correlation, and expected persistence of trade-cost shocks, not merely their size.

Looking forward, our results suggest that recent and prospective tariff increases applied to a broad set of U.S. trading partners could generate significantly larger pass-through than the targeted tariffs of 2018–2019, to the extent that breadth limits firms’ ability to substitute across source countries. Quantifying those effects as new data become available is a natural next step.

References

- AMITI, M., S. HEISE, AND A. WANG (2021): “High Import Prices along the Global Supply Chain Feed through to U.S. Domestic Prices,” Liberty street economics blog, Federal Reserve Bank of New York, november 8, 2021.
- BASLANDZE, S. AND S. FUCHS (2026): “The Price of Delay: Supply Chain Disruptions and Pricing Dynamics,” *Journal of International Economics*, 162, 104265.
- BLANCO, A., C. BOAR, C. J. JONES, AND V. MIDRIGAN (2024a): “The Inflation Accelerator,” Working paper, National Bureau of Economic Research.
- (2024b): “Nonlinear Inflation Dynamics in Menu Cost Economies,” Working paper, National Bureau of Economic Research.
- BONADIO, B., Z. HUO, A. A. LEVCHENKO, AND N. PANDALAI-NAYAR (2021): “Global Supply Chains in the Pandemic,” *Journal of International Economics*, 133, 103534.
- BORUSYAK, K., P. HULL, AND X. JARAVEL (2022): “Quasi-Experimental Shift-Share Research Designs,” *Review of Economic Studies*, 89, 181–213.
- CALIENDO, L. AND F. PARRO (2015): “Estimates of the Trade and Welfare Effects of NAFTA,” *Review of Economic Studies*, 82, 1–44.
- CAMPA, J. M. AND L. S. GOLDBERG (2005): “Exchange Rate Pass-Through into Import Prices,” *Review of Economics and Statistics*, 87, 679–690.
- CARRIÈRE-SWALLOW, Y., P. DEB, D. FURCERI, D. JIMÉNEZ, AND J. D. OSTRY (2023): “Shipping Costs and Inflation,” *Journal of International Money and Finance*, 130, 102771.

- CAVALLO, A. AND O. KRYVTSOV (2023): “What Can Stockouts Tell Us About Inflation? Evidence from Online Micro Data,” *Journal of International Economics*, 146, 103769.
- CUBA-BORDA, P., A. QUERALTO, R. REYES-HEROLES, AND M. SCARAMUCCI (2025): “Trade Costs and Inflation Dynamics,” Working Paper 2508, Federal Reserve Bank of Dallas, revised June 2025.
- DI GIOVANNI, J., Ş. KARAKAŞ, A. OZHAN, AND A. SALA (2022): “Global Supply Chain Pressures, International Trade, and Inflation,” NBER Working Paper 30240, National Bureau of Economic Research.
- GAGLIARDONE, L., M. GERTLER, S. LENZU, AND J. TIELENS (2025): “Micro and Macro Cost-Price Dynamics in Normal Times and During Inflation Surges,” Tech. rep., Working Paper, new York University and National Bank of Belgium.
- GILCHRIST, S., R. SCHOENLE, J. SIM, AND E. ZAKRAJŠEK (2017): “Inflation Dynamics During the Financial Crisis,” *American Economic Review*, 107, 785–823.
- GOLDBERG, P. K. AND M. M. KNETTER (1997): “Goods Prices and Exchange Rates: What Have We Learned?” *Journal of Economic Literature*, 35, 1243–1272.
- GOLDSMITH-PINKHAM, P., I. SORKIN, AND H. SWIFT (2020): “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, 110, 2586–2624.
- GOLOSOV, M. AND R. E. J. LUCAS (2007): “Menu Costs and Phillips Curves,” *Journal of Political Economy*, 115, 171–199.
- GOPINATH, G., O. ITSKHOKI, AND R. RIGOBON (2010): “Currency Choice and Exchange Rate Pass-Through,” *American Economic Review*, 100, 304–336.

- HUMMELS, D. (2007): “Transportation Costs and International Trade in the Second Era of Globalization,” *Journal of Economic Perspectives*, 21, 131–154.
- HUMMELS, D., V. LUGOVSKYY, AND A. SKIBA (2009): “The Trade Reducing Effects of Market Power in International Shipping,” *Journal of Development Economics*, 89, 84–97.
- JIAO, Y., T. LAN, Y. LIU, AND X. ZHOU (2026): “From Ports to Prices: The Inflationary Effects of Global Supply Chain Disruptions,” *AEA Papers and Proceedings*, 116, 646–651.
- KAMAL, F. AND W. OUYANG (2020): “Identifying U.S. Merchandise Traders: Integrating Customs Transactions with Business Administrative Data,” CES Working Paper CES-20-28, Center for Economic Studies, U.S. Census Bureau.
- LAFROGNE-JOUSSIER, R., J. MARTIN, AND I. MÉJEAN (2023): “Supply Shocks in Supply Chains: Evidence from the Early Lockdown in China,” *IMF Economic Review*, 71, 170–215.
- MICHAIL, N. A., K. D. MELAS, AND L. CLEANTHOUS (2022): “The Relationship Between Shipping Freight Rates and Inflation in the Euro Area,” *International Economics*, 172, 40–49.
- MIDRIGAN, M. (2011): “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, 79(4), 1139–1180.
- NAKAMURA, E. AND J. STEINSSON (2008): “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 123, 1415–1464.
- SHAPIRO, J. S. (2016): “Trade Costs, CO₂, and the Environment,” *American Economic Journal: Economic Policy*, 8, 220–254.

Appendices

A Merging PPI and LFTTD Databases

We merge the PPI and LFTTD databases using a multi-step name-matching algorithm designed to maximize match rates while minimizing false positives. Because the two databases do not share a common firm identifier, we rely on business names as the primary matching variable.

The matching procedure proceeds in several stages. First, we conduct an iterative word match routine that is repeated on two versions of the business name: (i) the raw original text from the PPI records and (ii) a cleaned version of the name (converted to uppercase, with punctuation such as “&” and “-” removed, and common business suffixes such as “INC,” “CO,” “LLC,” “CORP,” and “LTD” standardized and dropped).

The word match routine is carried out in three steps. In the first step, we directly match the full cleaned name string against the LFTTD database. This step captures exact or near-exact name matches. In the second step, the name string is tokenized into its constituent words, and we attempt to match the first five words of the name against the LFTTD. This step captures cases where business names differ only in trailing words (e.g., division names or location identifiers). The third step is identical to the second step except that we use a maximum of four words only, capturing additional matches where names share a common root but differ in specificity.

After exhausting the deterministic matching steps, we apply a fuzzy name matching procedure to the remaining unmatched records. Specifically, we use the SAS DQ-MATCH function with a sensitivity parameter of 95 to generate a standardized version of the cleaned business name. We retain matches to a unique firm identifier. For cases in which a PPI business name matches to multiple LFTTD firm identifiers, we resolve the ambiguity by matching the name strings word by word and retaining only those

matches for which at least two constituent words agree, starting from the maximum possible number of matched words.

To further validate the matches, we impose several consistency checks. We require that matched firms belong to the same broad industry group (at the 2-digit NAICS level) and, where geographic information is available, that they are located in the same state. These filters help eliminate spurious matches arising from common business names (e.g., “American Manufacturing” or “National Products”).

The resulting matched sample contains approximately 6,000 unique manufacturing firms that appear in both the PPI and LFTTD databases during the 2005–2022 period. These firms account for a substantial share of total U.S. manufacturing imports, suggesting that our sample is representative of the population of importing manufacturers.

B Robustness of the Bartik Instrument

In this appendix, we provide additional evidence on the validity and robustness of our Bartik-style instrumental variables strategy.

A potential concern with shift-share instruments is that the “shares” (in our case, the firm’s predetermined trade route exposure $z_{j,r,t}$) may be correlated with unobserved firm characteristics that independently affect producer prices. To address this concern, we conduct several robustness exercises.

First, we verify that the first-stage relationship between the instrument $B_{j,t}$ and the endogenous variable $\ln c_{j,t}$ is strong. Across all specifications, the first-stage F-statistic exceeds conventional thresholds for weak instrument concerns (ranging from 25 to over 100, depending on the specification), confirming that the instruments are relevant.

Second, we examine the sensitivity of our results to alternative definitions of the trade route exposure shares. In our baseline specification, we use a 60-month rolling window to compute $z_{j,r,t}$. As a robustness check, we re-estimate the model using 36-

month and 84-month windows. The results are quantitatively similar across all windows, suggesting that our findings are not driven by the specific choice of lag structure.

Lastly, we explore alternative leave-out strategies for constructing the route-level shipping cost component $C_{r,t}$. In our baseline, we exclude all firms in the same 2-digit NAICS industry. As alternatives, we exclude firms at the 3-digit and 4-digit NAICS levels. The point estimates remain stable across these specifications, providing further confidence in the exogeneity of our instruments.

C Value Functions

Three value functions describe the problem: the value of not adjusting, $\mathcal{V}_t^N$; the value of adjusting before the menu cost is paid, $\mathcal{V}_t^A$; and the value at the start of the period, before the menu-cost draw, $\mathcal{V}_t$:

$$\mathcal{V}_t^N \left(\frac{P_{t-1}(f)}{P_t}, z_t(f) \right) = \Pi_t \left(\frac{P_{t-1}(f)}{P_t}, z_t(f) \right) + \beta E_t \mathcal{V}_{t+1} \left(\frac{P_{t-1}(f)}{P_{t+1}}, z_{t+1}(f) \right), \quad (39)$$

$$\mathcal{V}_t^A(z_t(f)) = \max_{P_t(f)} \left\{ \Pi_t \left(\frac{P_t(f)}{P_t}, z_t(f) \right) + \beta E_t \mathcal{V}_{t+1} \left(\frac{P_t(f)}{P_{t+1}}, z_{t+1}(f) \right) \right\}, \quad (40)$$

$$\mathcal{V}_t \left(\frac{P_{t-1}(f)}{P_t}, z_t(f) \right) = E_{\chi_t(f)} \left[\max_{I_t(f) \in \{0,1\}} \left\{ I_t(f) [\mathcal{V}_t^A(z_t(f)) - \chi_t(f)] + (1 - I_t(f)) \mathcal{V}_t^N \left(\frac{P_{t-1}(f)}{P_t}, z_t(f) \right) \right\} \right], \quad (41)$$

where E_t is taken over $(\eta_{t+1}, \tilde{\epsilon}_{t+1}(f))$ and $E_{\chi_t(f)}$ over the menu-cost draw, and the flow profit in (39) is evaluated at $P_{t-1}(f)$ because a firm that does not adjust charges $P_t(f) = P_{t-1}(f)$, so that $\tilde{p}_t(f) = -x_{t-1}(f)$; the firm adjusts if and only if $\chi_t(f) \leq \mathcal{V}_t^A(z_t(f)) - \mathcal{V}_t^N(P_{t-1}(f)/P_t, z_t(f))$. The value of adjusting does not depend on the

inherited price: once the firm resets, $P_{t-1}(f)$ enters neither the flow profit nor the continuation. The time subscripts carry the dependence on the deterministic crisis path; in normal times the three value functions are stationary. The profit-scale ratio implied by the law of motion of $z_t(f)$ is

$$\frac{\bar{\Pi}_{t+1}(z_{t+1}(f))}{\bar{\Pi}_t(z_t(f))} = e^{(\sigma-1)\Delta z_{t+1}(f)} = e^{(\sigma-1)(\tilde{c}_{t+1}(f) - \phi s(f)\Delta u_{t+1})}, \quad (42)$$

with conditional mean $E_t \left[\frac{\bar{\Pi}_{t+1}(z_{t+1}(f))}{\bar{\Pi}_t(z_t(f))} \right] = e^{\frac{(\sigma-1)^2}{2}\sigma_\epsilon^2} e^{-(\sigma-1)\phi s(f)\Delta u_{t+1}}$. Hereafter, we assume that its conditional mean is approximately one.

Following GGLT, we scale the three value functions by the profit scale $\bar{\Pi}_t(z_t(f))$, using $\chi_t(f)/\bar{\Pi}_t(z_t(f)) = \tilde{\chi}_t(f)$:

$$V_t(x_{t-1}(f)) \equiv \frac{\mathcal{V}_t^N}{\bar{\Pi}_t} - \frac{1}{1-\beta}, \quad V_t^a \equiv \frac{\mathcal{V}_t^A}{\bar{\Pi}_t} - \frac{1}{1-\beta}, \quad \mathcal{E}_t(x_{t-1}(f)) \equiv \frac{\mathcal{V}_t}{\bar{\Pi}_t} - \frac{1}{1-\beta}.$$

D Proofs

Lemma 1 (Quadratic loss). *A second-order Taylor expansion of the profit function (19) around $\tilde{p}_t(f) = 0$ gives*

$$\Pi_t \left(\frac{P_t(f)}{P_t}, z_t(f) \right) = \bar{\Pi}_t(z_t(f)) \left(1 - \frac{\sigma(\sigma-1)}{2} \tilde{p}_t(f)^2 \right) + O(\tilde{p}_t(f)^3), \quad (43)$$

where $\tilde{p}_t(f) = p_t(f) - p_t^o(f)$ and $\bar{\Pi}_t(z_t(f)) \equiv \frac{Y}{\sigma} e^{(\sigma-1)z_t(f)}$ is the gap-zero profit level.

Proof of Lemma 1. Expand $F(\tilde{p}) \equiv e^{(1-\sigma)\tilde{p}} - e^{-\bar{\mu}-\sigma\tilde{p}}$: $F(0) = 1 - e^{-\bar{\mu}} = 1/\sigma$ (since $e^{-\bar{\mu}} = \frac{\sigma-1}{\sigma}$); $F'(0) = (1-\sigma) + \sigma e^{-\bar{\mu}} = (1-\sigma) + (\sigma-1) = 0$ (the static optimum); $F''(0) = (1-\sigma)^2 - \sigma^2 e^{-\bar{\mu}} = (\sigma-1)^2 - \sigma(\sigma-1) = -(\sigma-1)$. Hence $F(\tilde{p}) = \frac{1}{\sigma} - \frac{\sigma-1}{2} \tilde{p}^2 + O(\tilde{p}^3)$, and multiplying by $Y e^{(\sigma-1)z}$ gives (20). $\square$

Lemma 2 (Expected-value form of the adjustment decision). *Let $V_t^a \equiv \max_x V_t(x)$ be the value at the optimal reset gap and $G_t(x) \equiv V_t^a - V_t(x) \geq 0$ the gain from adjusting for a firm with ex-ante gap x . The expected value of entering the period with gap x , before the menu-cost draw and under the optimal adjustment policy, is*

$$\mathcal{E}_t(x) \equiv E_{\tilde{\chi}}[\max\{V_t^a - \tilde{\chi}, V_t(x)\}] = \begin{cases} V_t(x) + \frac{G_t(x)^2}{2\bar{\chi}}, & G_t(x) \leq \bar{\chi}, \\ V_t^a - \frac{\bar{\chi}}{2}, & G_t(x) > \bar{\chi}, \end{cases} \quad (44)$$

and the probability of adjustment—the generalized hazard function (GHF)—is

$$h_t(x) = \Pr(\tilde{\chi} \leq G_t(x)) = \min\left\{\frac{G_t(x)}{\bar{\chi}}, 1\right\} = -\frac{V''(0)}{2\bar{\chi}}x^2 + o(x^2), \quad (45)$$

the last expression being the second-order expansion around the reset point in the stationary case.

Proof of Lemma 2. The firm adjusts when the draw satisfies $\tilde{\chi} \leq G_t(x)$, paying the realized cost; hence for $G \leq \bar{\chi}$,

$$\mathcal{E}_t(x) = \int_0^G (V_t^a - c) \frac{dc}{\bar{\chi}} + \left(1 - \frac{G}{\bar{\chi}}\right) V_t(x) = V_t(x) + \frac{G^2}{2\bar{\chi}},$$

and for $G > \bar{\chi}$ the firm always adjusts, paying $E[\tilde{\chi}] = \bar{\chi}/2$. The hazard is the uniform CDF evaluated at G ; the quadratic expansion follows from $G_t(x) = -\frac{1}{2}V''(0)x^2 + o(x^2)$ when the reset point is at the maximum of a smooth V . $\square$

Lemma 3 (Survival-weighted reset price). *Consider a firm that adjusts at date t . Let $\mathcal{S}_{t,s} \equiv \prod_{\tau=1}^s (1 - h_{t+\tau}(x_{t+\tau}))$, with $\mathcal{S}_{t,0} \equiv 1$, be the probability that the new price is still in force at $t + s$ along the no-adjustment path—a random variable, since the hazards are evaluated at gaps that depend on the realized innovations $(\tilde{\epsilon}, \eta)$. Then the optimal*

reset price satisfies

$$p_t^*(f) = \frac{E_t \left[\sum_{s \geq 0} \beta^s \mathcal{S}_{t,s} p_{t+s}^o(f) \right]}{E_t \left[\sum_{s \geq 0} \beta^s \mathcal{S}_{t,s} \right]}. \quad (46)$$

Writing $S_{t,s} \equiv E_t \mathcal{S}_{t,s}$ for the expected survival probability and $\omega_{t,s} \equiv \beta^s S_{t,s} / \sum_{r \geq 0} \beta^r S_{t,r}$, the reset price is the survival-discount weighted average of expected future targets,

$$p_t^*(f) = \sum_{s \geq 0} \omega_{t,s} E_t \left[p_{t+s}^o(f) \right], \quad (47)$$

up to the covariance terms $\sum_s \beta^s \text{Cov}_t(\mathcal{S}_{t,s}, p_{t+s}^o(f)) / \sum_r \beta^r S_{t,r}$, which vanish to first order in normal times (the survival product is an even, and the target innovation an odd, function of the symmetric innovations around the zero reset gap) and are of second order in the product of the window drift and the innovation variance.

Proof of Lemma 3. Since $V_t^a = \max_x V_t(x)$, the reset gap satisfies the first-order condition $V_t'(x_t^*) = 0$. Differentiating (22) and using the envelope property of the adjustment decision, $\mathcal{E}'_{t+1}(x) = (1 - h_{t+1}(x)) V'_{t+1}(x)$ (from Lemma 2: at the threshold draw $\tilde{\chi} = G_{t+1}$ the firm is indifferent, so shifting the threshold has no first-order value effect), gives

$$\sigma(\sigma - 1) x_t^* = \beta E_t \left[(1 - h_{t+1}) V'_{t+1}(x_t) \right],$$

and

$$V'_{t+1}(x_t) = -\sigma(\sigma - 1) x_t + \beta E_{t+1} \left[(1 - h_{t+2}) V'_{t+2}(x_{t+1}) \right].$$

The second equation being the envelope condition for the value of not adjusting along the no-adjustment path. Substituting the second into the first repeatedly and using transversality yields $E_t \left[\sum_{s \geq 0} \beta^s \mathcal{S}_{t,s} (p_{t+s}^o(f) - p_t^*(f)) \right] = 0$, since the gap of the date- t price at $t + s$ along that path is $p_{t+s}^o(f) - p_t^*(f)$. As $p_t^*(f)$ is known at t it factors out, giving (46). The hazard path is held fixed without loss: by the envelope property its dependence on the reset price contributes zero at the optimum.

The only approximation in the deterministic-weight form (47) is the survival–target covariance stated in the lemma. Equivalently, (46) is the first-order condition of $\min_p E_t \sum_{s \geq 0} \beta^s \mathcal{S}_{t,s} (p - p_{t+s}^o(f))^2$ with the survival weights taken as given. $\square$

Proof of Proposition 1. Decompose $E_t [p_{t+h}(f)]$ by the date of the *last* reset in $(t, t+h]$. Under the constant hazard the probability of no reset is θ^h , and the probability that the last reset falls at $t+j$ is $(1-\theta)\theta^{h-j}$ (adjust at $t+j$, survive the remaining $h-j$ months); the weights sum to one. Hence

$$E_t [p_{t+h}(f)] = \theta^h E_t [p_t(f)] + \sum_{j=1}^h (1-\theta)\theta^{h-j} E_t [p_{t+j}^*(f)]. \quad (48)$$

Two exposure derivatives complete the calculation. First, a price *reset* at window date a builds in the survival-discount forward average of the remaining path (Lemma 3):

$$p_a^*(f) = \sum_{s \geq 0} \omega_s E_a [p_{a+s}^o(f)], \quad (49)$$

so that,

$$\frac{\partial E_t [p_a^*(f)]}{\partial \widehat{\tau}_t(f)} = \phi s(f) \sum_{s \geq 0} \omega_s u_{a+s} \equiv \phi s(f) \varphi(a), \quad (50)$$

the forward-looking crisis loading. Second, the *inherited* price at t was itself set at some past reset date; averaging over the (geometric) distribution of that date and applying (50) at each,

$$\frac{\partial E_t [p_t(f)]}{\partial \widehat{\tau}_t(f)} = \phi s(f) \ell_t, \quad \ell_\tau \equiv (1-\theta) \sum_{j \geq 0} \theta^j \varphi(\tau-j) \mathbf{1}\{\tau-j \geq 1\}, \quad (51)$$

the probability-weighted average of the loadings embedded at past reset dates ($\ell_\tau = 0$ for $\tau \leq 0$: pre-window prices carry none of the shock).

Differentiating (48) with respect to $\widehat{\tau}_t(f)$ using (50)–(51) gives the first equality of (30). Next by splitting the sum in (51) at horizon $\tau = t + h$ into resets after t (the terms $j < h$, giving $(1 - \theta) \theta^{h-j} \varphi(t + j)$ after reindexing) and resetting at or before t (whose total mass is θ^h times the same weights that define ℓ_t), we obtain:

$$\ell_{t+h} = \theta^h \ell_t + \sum_{j=1}^h (1 - \theta) \theta^{h-j} \varphi(t + j). \quad (52)$$

□

Proof of Corollary 1. First, substitute $\bar{u}_{a+s} = \rho^s \bar{u}_a$ into (50): $\varphi(a) = \bar{u}_a \sum_{s \geq 0} \omega_s \rho^s = \lambda(\rho) \bar{u}_a$ by the geometric sum.

Second, substitute (1) into (51) and use the two-rate convolution identity $\sum_{j=0}^{\tau-1} \theta^j \rho^{\tau-1-j} = (\rho^\tau - \theta^\tau) / (\rho - \theta)$; treating τ as continuous, $d\ell/d\tau \propto \rho^\tau \ln \rho - \theta^\tau \ln \theta$, positive at $\tau = 0$ (it equals $\ln(\rho/\theta) > 0$) and eventually negative, with a single sign change at τ^* since $(\theta/\rho)^\tau$ is monotone.

Lastly, to show the monotonicity in t , it is sufficient to note that the coefficient is proportional to $\rho^{h+1} [1 - (\theta/\rho)^{t+h}]$, increasing in t because $\theta < \rho$; letting $t \rightarrow \infty$ gives (36). □