The Impact of Shipping Costs on U.S. Producer Prices: Micro-Level Evidence from the Pandemic

Simon Gilchrist* Bin Wei† Vivian Z. Yue‡ Egon Zakrajšek§

March 2026

Abstract

This paper examines how fluctuations in international shipping costs pass through to U.S. producer prices. Using confidential shipment-level trade transactions from the Census Bureau, we construct time-varying firm-level measures of ad valorem shipping costs and then merge these costs with confidential monthly firm-level producer prices from the Bureau of Labor Statistics for about 6,000 U.S. manufacturing importing firms over the 2005–2022 period. Using both OLS and IV techniques—with the latter based on Bartik-style instruments that exploit exogenous variation in route-level shipping costs—we find that the pass-through of shipping costs to producer prices is negligible during "normal" times, but increased substantially during the COVID-19 pandemic, accounting for 17 percent of the 12-month PPI inflation at its peak. We develop a theoretical framework that rationalizes these findings by embedding multi-country sourcing into a menu-cost pricing model with variable markups. Calibrated to match the data, the model reproduces the near-zero normal-times pass-through and the crisis pass-through of approximately 3.6 percent at the 12-month horizon, and also provides a quantitative decomposition of the pass-through.

JEL Classification: E31, F14, F41, L11
Keywords: supply chain disruption, shipping costs, pass-through, domestic producer prices, inflation

The views expressed herein are solely those of the authors and should not be reported as representing the views of the Federal Reserve Bank of Atlanta, the Federal Reserve Bank of Boston, the principals of the Board of Governors, or the Federal Reserve System.

*New York University and NBER. Email: sg40@nyu.edu
†Federal Reserve Bank of Atlanta. Email: bin.wei@atl.frb.org
‡Emory University, NBER, and CEPR. Email: vyue@emory.edu
§Federal Reserve Bank of Boston and CEPR. Email: egon.zakrajsek@bos.frb.org

1 Introduction

Understanding how international trade cost shocks 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. firms. 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—at a granular level—firm-level trade costs and firm-level output prices.

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 provide firm-level monthly output prices. By matching these two datasets 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. 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 identification strategy addresses a key endogeneity concern: fluctuations in international shipping costs may be correlated with changes in demand conditions that independently affect producers' domestic prices. To isolate exogenous supply-side variation in shipping 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. This so-called leave-one-out construction mitigates the reflection problem inherent in shift-share designs.

Our empirical analysis yields two main findings. During "normal" times, the pass-through of shipping costs to producer prices is negligible. Panel OLS regressions of future producer prices on current ad valorem shipping costs—controlling for industry-by-month fixed effects, firm fixed effects, and lagged prices—yield coefficients that are economically negligible and statistically insignificant at horizons three to 12 months. This finding is consistent with the view that in tranquil periods, firms can adjust along several margins—substituting suppliers, absorbing cost increases into markups, or renegotiating contracts—to insulate their domestic output prices from fluctuations in international shipping costs.

During the COVID-19 pandemic, by contrast, when unprecedented disruptions in global supply chains sent international shipping costs soaring, 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[pre-determined 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. At the 12-month horizon, our IV estimates imply that shipping cost increases during the pandemic accounted for approximately 17 percent of the 12-month manufacturing PPI inflation. The fact that the IV estimates are substantially larger than the corresponding OLS estimates is consistent with the downward bias in OLS estimates arising from demand shocks that simultaneously reduce both import volumes (and hence measured shipping costs) and domestic producer prices.

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 due to menu costs. To illustrate the mechanism, we develop a theoretical framework that embeds multi-country sourcing into a dynamic pricing model with menu costs and variable markups à la Atkeson and Burstein (2008). In the model, a firm's pass-through of shipping cost shocks to its output price decomposes into the product of three components: (i) a direct cost channel, whose magnitude depends on the import expenditure share and on whether the shock is idiosyncratic to one supplier or common to all; (ii) a markup absorption channel, whereby strategic complementarity in pricing causes the firm to compress its markup in response to a cost increase, attenuating pass-through by a factor of $(1-\Gamma)$; and (iii) a menu-cost friction $\mathcal{F}$, which we show is increasing and convex in the size of the shock relative to the firm's inaction band. The interaction of supplier diversification and nonlinear price adjustment generates the state-dependent pass-through documented in the data. In normal times, route-specific shipping cost shocks are diversified across N suppliers, producing a small shift in the firm's optimal price that falls well within the inaction region; the resulting pass-through is negligible. During a crisis such as the pandemic, correlated shocks eliminate the diversification effect, multiplying the effective cost impact by a factor of N and pushing a substantial fraction of firms outside their inaction regions. A calibrated version of the model quantitatively matches the near-zero normal-times pass-through and the amplified crisis pass-through of approximately 3.6 percent at the 12-month horizon observed in the data.

The remainder of the paper is organized as follows. Section 2 discusses the related literature. Section 3 describes our data sources, measurement approach, and main empirical results. Section 4 outlines a theoretical framework that rationalizes our empirical findings. In Section 5 we presents calibration results and provides a quantitative decomposition of the pass-through. Section 6 concludes.

2 Related Literature

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

Trade cost pass-through: A large literature studies how trade costs transmit to domestic prices, with a central finding that pass-through is typically incomplete. Early work on exchange rate pass-through established this regularity for import prices (Goldberg and Knetter, 1997; Campa and Goldberg, 2005), while Gopinath, Itskhoki, and Rigobon (2010) show that the currency of invoicing plays a critical role. The pass-through of shipping costs has received less attention due to data limitations. Hummels, Lugovskyy, and Skiba (2009) find that freight cost pass-through to import prices varies across products, and Hummels (2007) provides comprehensive evidence on the magnitude of trade costs. Our contribution is to link shipping costs directly to firm-level producer prices, rather than to import prices or consumer prices, thereby capturing the pass-through at an earlier stage of the pricing chain.

Supply chain disruptions and pandemic-era inflation: A growing literature documents that pandemic-era supply chain disruptions contributed significantly to the inflation surge of 2021–2022. Amiti, Heise, and Weinstein (2023) find that the surge in container shipping rates significantly raised U.S. import prices, with effects that were larger and more persistent than in pre-pandemic periods. Di Giovanni et al. (2022) decompose U.S. inflation and find that supply-side factors accounted for a substantial share of the surge. At the macro level, Bai, FernandezVillaverde, LiZanetti (2025) construct a real-time port congestion index from satellite tracking data on container ships and, using a structural VAR identified through a search-friction model of goods markets, show that supply chain shocks were the primary driver of U.S. inflation in 2021. increases. At the micro level, Lafrogne-Joussier, Martin, and Méjean (2023) provide firm-level evidence from France showing that firms with greater exposure to disrupted supply chains were more likely to raise output prices. Our paper complements this literature by tracing the pass-through one step further along the pricing chain—to the producer prices of manufacturing firms that use imported inputs—and by documenting that pass-through is highly state-dependent.

Tariff incidence: A related literature examines who ultimately bears the burden of import duties. Amiti, Redding, and Weinstein (2019) find that the 2018–2019 U.S. tariffs on Chinese imports were almost entirely passed through to U.S. import prices. Fajgelbaum, Goldberg,

Kennedy, and Khandelwal (2020) reach a similar conclusion and estimate that the tariffs reduced U.S. real income by approximately $51 billion per year. Cavallo, Gopinath, Neiman, and Tang (2021) show that tariffs were fully reflected in border prices but only partially passed through to retail prices. Our paper extends this literature by examining pass-through at the producer price level, where we find that tariff increases were largely absorbed by manufacturers through markup compression.

Firm-level pricing and variable markups: Models of variable markups (Atkeson and Burstein, 2008; Kimball, 1995) predict that firms with market power partially absorb cost shocks by adjusting markups. Amiti, Itskhoki, and Konings (2019) and Berman, Martin, and Mayer (2012) provide direct evidence that larger firms exhibit lower pass-through, consistent with these models. Gilchrist, Schoenle, Sim, and Zakrajšek (2017) use PPI microdata—the same data source we employ—to study pricing dynamics during the Great Recession, and Nakamura and Steinsson (2008) document substantial heterogeneity in the frequency of price adjustment, with a median duration of 8–11 months for producer prices. This stickiness provides a natural explanation for the delayed and horizon-dependent pass-through that we document empirically.

3 Empirical Analysis

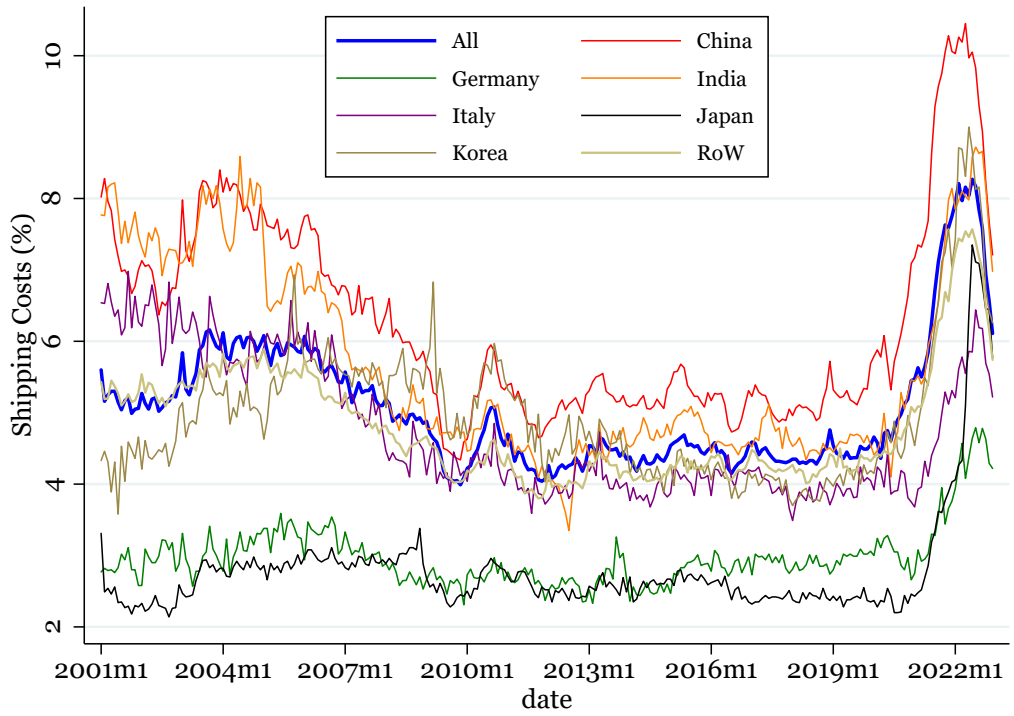
3.1 Data Sources and Methods

The main contribution of our paper is to construct firm-level trade cost and price change 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: 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 measures available in the LFTTD since the early 2000s, including shipment-level shipping costs (freight

FIGURE 1: Trade Costs



NOTE: This figure shows the value-weighted average of ad valorem shipping costs (freight and insurance charges) as a percentage of customs value for U.S. merchandise imports, disaggregated by major source country (China, Germany, India, Italy, Japan, Korea, and the rest of the world (RoW)). The thick blue line is the simple average across all source countries.
SOURCE: Authors' calculations using the U.S. Census Bureau's Longitudinal Foreign Trade Transactions Database.

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 Data: The Producer Price Index (PPI), published by the U.S. Bureau of Labor Statistics (BLS), 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, Schoenle, Sim, and Zakrajšek (2017), we use the micro data available from the BLS to construct firm-level inflation measures as the weighted average of price changes across goods within each firm:

$$\Delta p_{j,t} = \frac{1}{N_j} \sum_{i=1}^{N_j} w_{i,j,t} \Delta p_{i,j,t}, \tag{1}$$

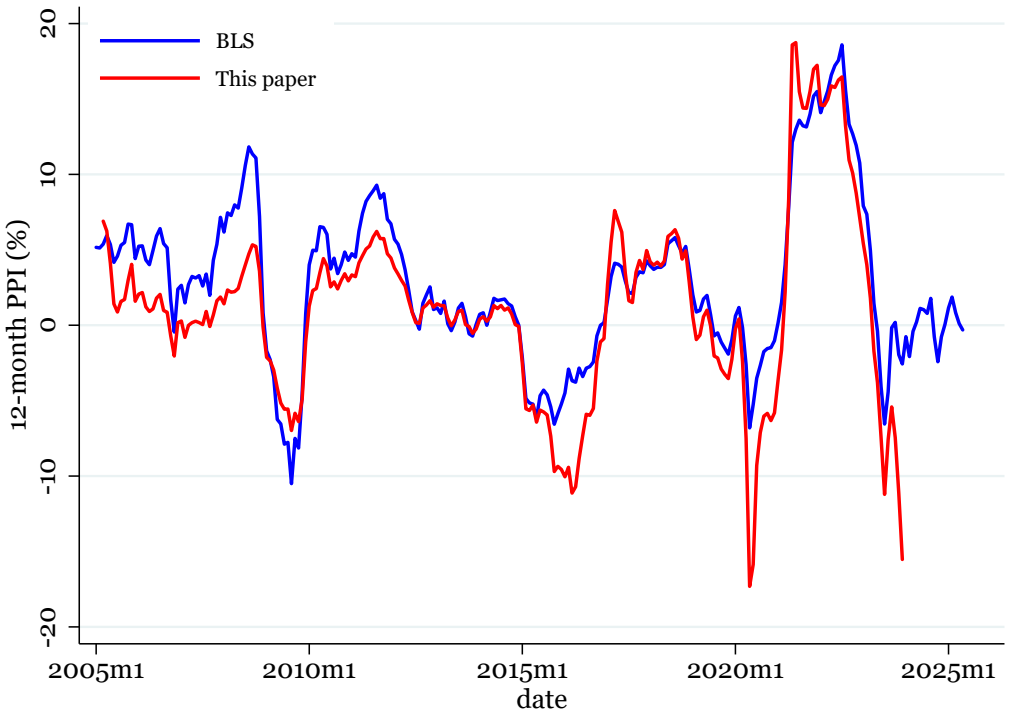
where $\Delta p_{i,j,t}$ denotes the monthly log-difference of price for each good i produced by firm j from month $t-1$ to month t and N_j is the number of goods reported by the firm to the BLS. The weights $w_{i,j,t}$ are based on the relative importance weights assigned by the BLS and on firm-level sales data.[1]

Figure 2 displays the 12-month PPI inflation rates. The blue line represents the 12-month log-difference of the official PPI index published by the BLS, while the red line shows the 12-month log-difference of prices in the manufacturing sector constructed using the BLS microdata. The two series track each other closely throughout the sample, providing validation of our measurement approach. 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.

The LFTTD–PPI 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.

[1] See Gilchrist, Schoenle, Sim, and Zakrajšek (2017) for a full description of how to construct firm-level price indices from the PPI data.

FIGURE 2: Producer Prices



NOTE: The blue line shows the 12-month log-difference of the official PPI, whereas the red line shows the weighted-average of the 12-month log-differences of firm-level prices for U.S. manufacturing firms in the PPI microdata; see equation (1) and text for details.
SOURCE: Authors' calculations using the Bureau of Labor Statistics' micro-level PPI database.

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 selected variables in the LFTTD–PPI dataset are presented in Table 1.

3.2 Empirical Specification

We now turn to estimating the pass-through of shipping costs to U.S. producer prices.To assess the impact of trade costs, we regress future producer prices on current shipping costs. We include interaction terms for crisis periods to capture possible differential effects—specifically, the COVID-19 pandemic for shipping costs and the initial phase of the U.S.-China trade war for tariffs.

TABLE 1: Selected Summary Statistics

	Mean	StdDev	P10	P50	P90
A. *Trade costs*					
Ad valorem shipping costs (%)	4.82	4.15	1.02	3.67	9.85
Ad valorem duties (%)	2.94	3.52	0.00	1.85	7.42
No. of source countries	5.3	4.1	1	4	11
No. of trade routes	7.8	5.9	2	6	16
B. *Producer prices*					
1-month price change (%)	0.12	2.84	−2.15	0.00	2.38
12-month price change (%)	1.45	7.62	−5.21	0.87	8.73
No. of products per firms	3.7	2.8	1	3	7

NOTE: This table reports summary statistics for selected variables in the merged LFTTD–PPI panel dataset, 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 using the U.S. Census Bureau's Longitudinal Foreign Trade Transactions Database and the Bureau of Labor Statistics' micro-level PPI database.

Our baseline empirical specification is as follows:

$$p_{i,j,t+h} = \beta_N c_{j,t} + \beta_C (c_{j,t} \times \text{CRISIS}_t) + \sum_{l=1}^{12} \gamma_l p_{i,j,t-l} + \text{FE} + \epsilon_{i,j,t+h}, \tag{2}$$

where $p_{i,j,t+h}$ denotes the log price of product i (produced by firm j) in month $t+h$ and $c_{j,t}$ represents the log of ad valorem shipping costs incurred by firm j in month t. The 0/1-indicator variable CRISIS_t equals one during crisis period—from March 2020 to December 2022— and zero otherwise. The interaction term $c_{j,t} \times \text{CRISIS}_t$ allows us to estimate the potentially non-linear pass-through of fluctuations in international shipping costs to domestic prices.

The coefficient β_N captures the average pass-through of shipping costs to producer prices during "normal" (non-crisis) times. The coefficient β_C captures the *incremental* pass-through during crisis periods, so that the total pass-through during crises is given by $\beta_N + \beta_C$. We estimate the model at multiple horizons ($h = 3, 6, 12$ months) to trace out the dynamics of the shipping 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.

3.3 Main Results

We begin by examining the pass-through of shipping costs to producer prices using OLS to estimate equation (2). The results of this exercise are reported in Table 2.

TABLE 2: The Pass-through of Shipping Costs to Producer Prices – OLS Estimates

	Horizon (h months)					
	$h = 3$	$h = 6$	$h = 12$	$h = 3$	$h = 6$	$h = 12$
Regressor	(1)	(2)	(3)	(4)	(5)	(6)
$c_{j,t}$	0.0001	0.0001	0.0003	−0.0001	−0.0002	−0.0004
	(0.0001)	(0.0002)	(0.0004)	(0.0002)	(0.0003)	(0.0004)
$c_{j,t} \times \text{CRISIS}_t$	0.0010	0.0016	0.0029	0.0006	0.0013	0.0024
	(0.0004)	(0.0007)	(0.0014)	(0.0005)	(0.0007)	(0.0014)
R^2	0.92	0.87	0.78	0.92	0.88	0.81
Firm FEs	N	N	N	Y	Y	Y
Obs.	947,000	907,000	821,000	947,000	907,000	821,000

NOTE: The dependent variable is $p_{i,j,t+h}$, the log of the price of item i (produced by firm j) in month $t + h$. Explanatory variables: $c_{j,t}$ = log of the ad valorem shipping costs incurred by firm j in month t; and CRISIS_t = 0/1-indicator that equals one between March 2020 and December 2022 and zero otherwise. All specifications include lags $1, 2, \ldots, 12$ of $p_{i,j,t}$ (not reported), 3-digit NAICS industry-by-month fixed effects, and are estimated by OLS. Asymptotic standard errors clustered at the firm level are reported in parentheses.

The baseline coefficient on shipping costs, β_N, is small and statistically insignificant in the OLS specifications. In columns (1) through (3), which do not include firm fixed effects, the estimated pass-through is close to zero at all horizons. When we include firm fixed effects (columns 4 through 6), the estimates of β_N 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 coefficient β_N on the interaction term $c_{j,t} \times \text{CRISIS}_t$, 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 a mere 0.029 percent increase in producer prices when firm fixed effects are not included (column 3) and a 0.024 percent increase when firm fixed effects are included (column 6).

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

$$B_{j,t} = \sum_{r} z_{j,r,t} C_{r,t}, \tag{3}$$

where

$$\begin{aligned} z_{j,r,t} &= \frac{\sum_{t' \in [t-60, t-1]} \text{value}_{j,r,t'}}{\sum_{r'} \sum_{t' \in [t-60, t-1]} \text{value}_{j,r',t'}}; \\ C_{r,t} &= \frac{\sum_{k \in \text{NAICS2}(j)} \text{cost}_{k,r,t}}{\sum_{k \in \text{NAICS2}(j)} \text{value}_{k,r,t}}. \end{aligned}$$

Here, $z_{j,r,t}$ denotes the share of firm j's total international trade that was routed through trade route r during the 60 months preceding month t and captures the firm's pre-determined exposure to route-level trade 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 by NAICS2(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.

The identifying assumption underlying this instrument 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.

TABLE 3: The Pass-through of Shipping Costs to Producer Prices – IV Estimates

	Horizon (h months)					
	$h=3$	$h=6$	$h=12$	$h=3$	$h=6$	$h=12$
Regressor	(1)	(2)	(3)	(4)	(5)	(6)
$c_{j,t}$	0.0040	0.0075	0.0136	0.0066	0.0130	0.0216
	(0.0009)	(0.0016)	(0.0029)	(0.0123)	(0.0197)	(0.0356)
$c_{j,t} \times \text{CRISIS}_t$	0.0115	0.0201	0.0368	0.0112	0.0208	0.0357
	(0.0043)	(0.0076)	(0.0144)	(0.0045)	(0.0080)	(0.0144)
Firm FEs	N	N	N	Y	Y	Y
Obs.	947,000	907,000	821,000	947,000	907,000	821,000

NOTE: The dependent variable is $p_{i,j,t+h}$, the log of the price of item i (produced by firm j) in month $t+h$. Explanatory variables: $c_{j,t}$ = log of the ad valorem shipping costs incurred by firm j in month t; and CRISIS_t = 0/1-indicator that equals one between March 2020 and December 2022 and zero otherwise. The explanatory variables are instrumented by $B_{j,t}$ and $B_{j,t} \times \text{CRISIS}_t$, where $B_{j,t}$ is a Bartik-style intrument constructed using firm j's historical exposure to different trade routes interacted with route-level national average shipping costs (see the text for details). All specifications include lags $1, 2, \ldots, 12$ of $p_{i,j,t}$ (not reported), 3-digit NAICS industry-by-month fixed effects, and are estimated by 2SLS. Asymptotic standard errors clustered at the firm level are reported in parentheses.

Table 3 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 reported in Table 2. The baseline pass-through outside the crisis period (β_N) is positive and statistically significant when firm fixed effects are not included (columns 1 through 3). However, once firm effects are included (columns 4 through 6), 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) through (3) 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, import intensity, 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 coefficient on the crisis interaction term ($\beta_{scriptscriptstyleC}$) remains strongly statistically significant and is an order of magnitude larger than the OLS estimates. Its magnitude increases with the horizon, rising from approximately 0.01 at the three-month horizon to nearly 0.04 at the 12-month horizon, 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, representing an approximate 75 percent increase. Applying our IV estimate of $\beta_C \approx 0.036$ at the 12-month horizon, this implies a direct contribution of shipping cost increases to PPI inflation of approximately $0.036 \times \ln(1.75) \approx 2.0$ percentage points, or roughly 17 percent of the peak 12-month PPI inflation of approximately 12 percentage points for manufacturing firms.

Overall, the results point to limited pass-through of international shipping costs to producer prices in normal times, but significantly stronger effects during periods of acute supply chain disruption such as the COVID-19 crisis. 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 model in which firms can substitute across suppliers and trade routes to offset idiosyncratic cost shocks, but face binding constraints when cost increases are broad-based and highly correlated across all sourcing options. During the pandemic, shipping costs rose simultaneously across virtually all trade routes and origin countries, leaving firms with fewer avenues for adjustment. In this environment, firms were forced to either absorb the costs—compressing already-strained margins—or pass them through to output prices.

4 Model

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. The economy consists of a continuum of intermediate goods producers (firms) indexed by $i \in [0, 1]$, each producing a differentiated variety. There are N foreign countries indexed by $j = 1, \ldots, N$ from which firms can source intermediate inputs. The domestic economy is small relative to the world, so foreign prices are taken as given.

4.1 Production Technology

Each firm i produces output $Y_{i,t}$ using labor $L_{i,t}$ and a composite intermediate input $X_{i,t}$ according to a Cobb-Douglas technology:

$$Y_{i,t} = \Omega_{i,t} \, X_{i,t}^{\phi} \, L_{i,t}^{1-\phi}, \tag{4}$$

where $\Omega_{i,t}$ is firm-specific productivity and $\phi \in (0, 1)$ is the intermediate input share. The composite intermediate input aggregates N types of inputs, each associated with a foreign

country:

$$X_{i,t} = \prod_{j=1}^{N} X_{i,j,t}^{\gamma_j}, \qquad \sum_{j=1}^{N} \gamma_j = 1, \tag{5}$$

where γ_j represents the importance of type-j intermediate input in production.

Each type-j input combines domestic and foreign varieties that are imperfect substitutes via a CES aggregator:

$$X_{i,j,t} = \left[Z_{i,j,t}^{\zeta/(1+\zeta)} + a_j^{1/(1+\zeta)} M_{i,j,t}^{\zeta/(1+\zeta)} \right]^{(1+\zeta)/\zeta}, \tag{6}$$

where $Z_{i,j,t}$ and $M_{i,j,t}$ are quantities of domestic and imported varieties respectively, $a_j > 0$ is an efficiency parameter governing the productivity advantage of the foreign variety, and $\zeta > 0$ is the elasticity of substitution between domestic and foreign varieties.

4.2 Import Costs

The landed price of the foreign variety from country j in domestic currency is:

$$P_{j,t}^{M} = e_t \, U_{j,t} \, \tau_{j,t} \, (1 + d_{j,t}), \tag{7}$$

where e_t is the nominal exchange rate (domestic currency per unit of foreign currency), $U_{j,t}$ is the f.o.b. price in foreign currency, $\tau_{j,t} \geq 1$ is the ad valorem shipping cost (iceberg trade cost), and $d_{j,t} \geq 0$ is the ad valorem tariff rate. The total cost wedge between the foreign factory gate price and the domestic landed price is $\tau_{j,t}(1 + d_{j,t})$. We denote by $V_{j,t}^{*}$ the price of the domestic variety of type-j intermediate input.

Given output $Y_{i,t}$, the firm imports from all N countries and minimizes total variable costs subject to the production technology (4)–(6). We derive the total variable cost function in detail in Proposition 1.

Proposition 1 (Cost Minimization) *Given the Cobb-Douglas production function (4), the Cobb-Douglas aggregation of intermediate types (5), and the CES substitution between home and foreign varieties (6), the total variable cost function is given by*

$$TVC_t(Y_{i,t}) = \frac{C_t^{*}}{B_{N,t}^{\phi}} \cdot \frac{Y_{i,t}}{\Omega_{i,t}}, \tag{8}$$

where C_t^{*} is the unit cost index for a non-importing firm:

$$C_t^{*} = \left(\frac{V_t^{*}}{\phi} \right)^{\phi} \left(\frac{W_t}{1-\phi} \right)^{1-\phi}, \tag{9}$$

and $V_t^* \equiv \prod_{j=1}^{N}(V_{j,t}^*/\gamma_j)^{\gamma_j}$ *is the Cobb-Douglas price index over domestic input types, and* $B_{N,t}$ *is the cumulative cost-reduction factor from importing from all* N *countries:*

$$B_{N,t} = \prod_{j=1}^{N} b_{j,t}^{\gamma_j}, \tag{10}$$

and $b_{j,t}$ *is the cost-reduction factor from importing type-*j *intermediates:*

$$b_{j,t} = \left[1 + a_j \left(\frac{P_{j,t}^M}{V_{j,t}^*}\right)^{-\zeta}\right]^{1/\zeta}. \tag{11}$$

Proof. See Appendix A.

One immediate implication from the proposition above is that the marginal cost of a non-importing firm is $c_{0,t} = C_t^*/\Omega_{i,t}$, and importing from all N countries reduces marginal cost by a factor of $B_{N,t}^{\phi} > 1$.

4.3 Marginal Cost and Its Elasticity with Respect to Trade Costs

From Proposition 1, with all N suppliers always active, the firm's marginal cost is:

$$MC_{i,t} = \frac{C_t^*}{\Omega_{i,t} \cdot B_{N,t}^{\phi}}, \tag{12}$$

where $B_{N,t} = \prod_{j=1}^{N} b_{j,t}^{\gamma_j}$ is the cumulative cost-reduction factor from importing and $C_t^* = (V_t^*/\phi)^{\phi}(W_t/(1-\phi))^{1-\phi}$ is the common cost index. Note that $MC_{i,t}$ is deterministic conditional on prices and productivity—there is no expectation operator because the set of active suppliers is not stochastic.

The elasticity of marginal cost with respect to the shipping cost factor on imports from country j is

$$\frac{d\ln MC_{i,t}}{d\ln \tau_{j,t}} = \phi\,\gamma_j\,\tilde{a}_{j,t}, \tag{13}$$

where $\tilde{a}_{j,t}$ is the expenditure share of imports in type-j intermediate input:

$$\tilde{a}_{j,t} = \frac{a_j (P_{j,t}^M/V_{j,t}^*)^{-\zeta}}{1 + a_j (P_{j,t}^M/V_{j,t}^*)^{-\zeta}}. \tag{14}$$

In deriving equation (13), we have used the result $d\ln b_{j,t}/d\ln \tau_{j,t} = -\tilde{a}_{j,t}$, since the cost-reduction factor $b_{j,t}$ depends on the import price $P_{j,t}^M = e_t U_{j,t} \tau_{j,t}(1 + d_{j,t})$ and thus a change

in shipping costs $\tau_{j,t}$ affects marginal cost through $b_{j,t}$. The elasticity is negative because higher shipping costs raise the import price, reducing the gains from import variety.

For a common shipping cost shock affecting all countries proportionally:

$$\frac{d \ln MC_{i,t}}{d \ln \tau_t} = \phi \sum_{j=1}^{N} \gamma_j \, \tilde{a}_{j,t} = \phi \, s_M \equiv \alpha_d, \tag{15}$$

where $s_M \equiv \sum_{j=1}^{N} \gamma_j \tilde{a}_{j,t}$ is the aggregate import expenditure share. We collect these results as the next proposition.

Proposition 2 (Marginal Cost Elasticity) *Denote $\alpha_d \equiv \phi \, s_M$ and let $\omega_j = \gamma_j \tilde{a}_{j,t} / s_M$ be country j's share in the total import bill, then the elasticity of marginal cost with respect to the shipping cost factor on imports from country j is $\phi \, \gamma_j \, \tilde{a}_{j,t}$ or $\alpha_d \, \omega_j$. Moreover, the elasticity of marginal cost with respect to a common shipping cost shock is $\phi \, s_M$ or α_d.*[2]

4.4 Dynamic Pricing with Menu Costs

We now embed the cost structure developed so far into a dynamic pricing model with menu costs. This section develops the demand structure, formulates the firm's pricing problem, characterizes the optimal (S, s) policy, and describes the shock structure that distinguishes normal times from crisis periods.

Following Atkeson and Burstein (2008), we adopt a nested CES demand structure with two tiers of substitution: across sectors (elasticity η) and within sectors (elasticity $\rho > \eta$). The nested demand structure plays three essential roles in our model that a standard single-tier CES demand system cannot deliver. First, it generates *variable markups*. Under single-tier CES demand with a constant elasticity σ, every firm charges the same constant markup $\sigma/(\sigma-1)$ regardless of its size or the shocks it faces. Cost shocks are passed through to prices one-for-one (up to the constant markup), with no endogenous markup adjustment. The nested structure breaks this by making the effective demand elasticity $\sigma_{i,t} = \rho(1-S_{i,t}) + \eta S_{i,t}$ depend on the firm's market share $S_{i,t}$. Firms with larger market shares face less elastic demand and charge higher markups. When a cost shock raises a firm's price, the firm loses market share, faces more elastic demand, and its desired markup falls, partially absorbing the cost shock. This *strategic complementarity* is quantified by the parameter Γ in our pass-through decomposition.

Second, variable markups provide a distinct and empirically relevant *channel of incomplete pass-through* that operates independently of the menu-cost friction. Even if all firms

[2]For symmetric suppliers ($\gamma_j = 1/N$, $\tilde{a}_{j,t} = \tilde{a}$ for all j), this simplifies to α_d/N for a country-specific shock and $\alpha_d = \phi s_M$ for a common shock to all countries.

adjusted prices every period, markup absorption would still attenuate pass-through by a factor of $(1-\Gamma)$. This separation of the markup channel from the menu-cost channel is important for understanding which mechanism drives incomplete pass-through in which regime.

Third, the nested structure generates *heterogeneous pass-through across firms* based on their market shares. Large firms (high S) have more variable markups (higher Γ) and absorb a larger fraction of cost shocks through markup compression. Small firms (low S) behave closer to constant-markup firms. This heterogeneity is consistent with the empirical evidence in Amiti, Itskhoki, and Konings (2019) that large exporters exhibit lower pass-through due to variable markups.

Aggregation structure. The economy consists of a continuum of sectors indexed by $s \in [0,1]$. A competitive final good producer aggregates sectoral composites:

$$Y_t = \left[\int_0^1 Y_{s,t}^{(\eta-1)/\eta}\, ds\right]^{\eta/(\eta-1)}, \tag{16}$$

where $\eta > 1$ is the elasticity of substitution across sectors. Within each sector, a sectoral aggregator bundles differentiated varieties:

$$Y_{s,t} = \left[\int_0^1 \xi_{i,s}^{1/\rho}\, q_{i,s,t}^{(\rho-1)/\rho}\, di\right]^{\rho/(\rho-1)}, \tag{17}$$

where $\rho > \eta$ is the within-sector elasticity of substitution and $\xi_{i,s}$ is a demand shifter.

Similarly as in Atkeson and Burstein (2008), from the final good producer's expenditure minimization problem, we can solve for the sectoral demand function as

$$Y_{s,t} = \left(\frac{P_{s,t}}{P_t}\right)^{-\eta} Y_t, \tag{18}$$

where $P_t \equiv \left[\int_0^1 P_{s,t}^{1-\eta}\, ds\right]^{1/(1-\eta)}$ is the aggregate CES price index.

Within sector s, we can similarly express firm demand in terms of aggregate variables:

$$q_{i,s,t} = \xi_{i,s}\left(\frac{p_{i,s,t}}{P_{s,t}}\right)^{-\rho} \cdot \left(\frac{P_{s,t}}{P_t}\right)^{-\eta} Y_t = \xi_{i,s} \cdot p_{i,s,t}^{-\rho} \cdot P_{s,t}^{\rho-\eta} \cdot D_t,$$

where we define $D_t \equiv P_t^{\eta} Y_t$ as the aggregate demand index, and $P_{s,t} \equiv \left[\int_0^1 \xi_{i,s}\, p_{i,s,t}^{1-\rho}\, di\right]^{1/(1-\rho)}$ as the sectoral price index.

Market share and effective demand elasticity. The revenue share of firm i within sector s is:

$$S_{i,s,t} \equiv \frac{p_{i,s,t}\, q_{i,s,t}}{P_{s,t}\, Y_{s,t}} = \frac{\xi_{i,s}\, p_{i,s,t}^{1-\rho}\, P_{s,t}^{\rho}\, Y_{s,t}}{P_{s,t}\, Y_{s,t}} = \xi_{i,s} \left(\frac{p_{i,s,t}}{P_{s,t}} \right)^{1-\rho}. \tag{19}$$

Since $\rho > 1$, a higher relative price reduces market share.

The firm recognizes that changing $p_{i,s,t}$ affects the sectoral price index $P_{s,t}$:

$$\frac{\partial \ln P_{s,t}}{\partial \ln p_{i,s,t}} = \frac{\xi_{i,s}\, p_{i,s,t}^{1-\rho}}{P_{s,t}^{1-\rho}} = S_{i,s,t}.$$

That is, the firm's impact on the sectoral price index equals its market share. Taking the log of the demand function (19) and differentiating with respect to $\ln p_{i,s,t}$, we obtain the effective demand elasticity as follows:

$$\begin{aligned} \sigma_{i,s,t} &\equiv -\frac{\partial \ln q_{i,s,t}}{\partial \ln p_{i,s,t}} = \rho - (\rho - \eta) \cdot \frac{\partial \ln P_{s,t}}{\partial \ln p_{i,s,t}} = \rho - (\rho - \eta)\, S_{i,s,t} \\ &= \rho(1 - S_{i,s,t}) + \eta\, S_{i,s,t}. \end{aligned} \tag{20}$$

This is a weighted average of the within-sector elasticity ρ and the across-sector elasticity η, with weights $(1 - S_{i,s,t})$ and $S_{i,s,t}$. Since $\rho > \eta$, the effective elasticity is decreasing in market share: large firms face less elastic demand and optimally charge higher markups.

Desired markup and optimal price. A monopolistically competitive firm maximizes profit $(p - MC) \cdot q(p)$. The first-order condition yields the standard markup:

$$\mu_{i,s,t} = \frac{\sigma_{i,s,t}}{\sigma_{i,s,t} - 1}. \tag{21}$$

Since $\sigma_{i,s,t}$ depends on $S_{i,s,t}$, which depends on $p_{i,s,t}$, the markup is *endogenous*.

Hereafter, we omit the sectoral index s from the subscript to simplify notation. The desired (frictionless) optimal price is:

$$p_{i,t}^{*} = \mu_{i,t} \cdot MC_{i,t}. \tag{22}$$

When a cost shock hits, the desired price changes. The variable markup creates strategic complementarity: a firm that faces a cost increase wants to raise its price, but doing so reduces its market share, which raises its demand elasticity and compresses its desired markup, partially offsetting the cost increase.

4.4.1 Optimal price setting as a (S, s) Policy

Each firm faces the following menu cost problem in which there is a fixed cost $\kappa > 0$ (measured in units of steady-state revenue) to change its nominal price. The firm's state variable is the *price gap*:

$$x_{i,t} = \ln p_{i,t} - \ln p^*_{i,t}, \tag{23}$$

which measures the log deviation of the firm's actual price from its optimal price.

The key feature of menu-cost pricing is that the firm's nominal price $p_{i,t}$ remains fixed unless the firm pays κ to adjust. Conditional on the firm *not adjusting* between t and $t+1$ (i.e., $p_{i,t+1} = p_{i,t}$), we can derive the law of motion for the price gap as follows. Write the price gap at $t+1$:

$$x_{i,t+1} = \ln p_{i,t+1} - \ln p^*_{i,t+1}. \tag{24}$$

Since $p_{i,t+1} = p_{i,t}$ (no adjustment), substitute $\ln p_{i,t+1} = \ln p_{i,t}$:

$$x_{i,t+1} = \ln p_{i,t} - \ln p^*_{i,t+1}.$$

Adding and subtracting $\ln p^*_{i,t}$:

$$x_{i,t+1} = \underbrace{(\ln p_{i,t} - \ln p^*_{i,t})}_{= x_{i,t}} - \underbrace{(\ln p^*_{i,t+1} - \ln p^*_{i,t})}_{= \Delta \ln p^*_{i,t+1}}, \tag{25}$$

which yields:

$$x_{i,t+1} = x_{i,t} - \Delta \ln p^*_{i,t+1}, \tag{26}$$

where $\Delta \ln p^*_{i,t+1} = \ln p^*_{i,t+1} - \ln p^*_{i,t}$ is the log change in the optimal price between t and $t+1$, driven by shocks to marginal cost and desired markups. If the firm *does* adjust by paying κ and resetting its price to $p_{i,t+1} = p^*_{i,t+1}$, then $x_{i,t+1} = 0$. The complete law of motion is therefore:

$$x_{i,t+1} = \begin{cases} x_{i,t} - \Delta \ln p^*_{i,t+1} & \text{if no adjustment (inaction region),} \\ 0 & \text{if adjustment (price reset to } p^*\text{).} \end{cases} \tag{27}$$

The economic content is transparent: while the firm's nominal price is frozen, the optimal price keeps moving due to cost and demand shocks. Each period without adjustment, the price gap drifts by exactly $-\Delta \ln p^*_{i,t+1}$. If the optimal price rises ($\Delta \ln p^* > 0$, e.g., due to a cost increase), the gap x falls—the actual price is now "too low" relative to the optimum—pushing x toward the lower adjustment boundary s_L.

We now derive the flow profit function $\pi(x)$ that appears in the Bellman equation, and explain how it connects to the cost structure derived earlier. Recall that the firm charges

price $p_{i,t}$ and faces demand $q_{i,t}$ from (19). For the purpose of the menu-cost problem, we work with a local approximation around the steady state, treating the sectoral price index $P_{s,t}$ and aggregate variables as given (this is standard in the menu-cost literature; see Golosov and Lucas (2007)).

The firm's flow profit is given by $\Pi_{i,t} = p_{i,t} \cdot q_{i,t} - MC_{i,t} \cdot q_{i,t}$, where $MC_{i,t}$ is the marginal cost from the multi-country sourcing structure in equation (12). Normalize flow profit by steady-state revenue $p^* \cdot q(p^*)$. In Appendix A, we derive the firm's flow profit as follows:

$$\pi(x) = \frac{p_{i,t} \cdot q(p_{i,t}) - MC_{i,t} \cdot q(p_{i,t})}{p^* \cdot q(p^*)} = e^{(1-\theta)x} - \frac{\theta - 1}{\theta} e^{-\theta x}, \tag{28}$$

where $\theta \equiv \bar{\sigma}$ denotes the stead-state value of demand elasticity. Note that the first term is (normalized) revenue: it increases with x for small positive x because the firm earns a higher margin, but eventually decreases because demand falls off. The second term is (normalized) cost: it decreases with x because higher prices reduce quantity and hence production costs.

The firm's dynamic programming problem is:

$$V(x) = \pi(x) + \beta \, \mathbb{E}\left[\max\left\{V(x'),\ \sup_{x''} V(x'') - \kappa\right\}\right], \tag{29}$$

where $x' = x - \Delta \ln p^*$ is the price gap in the next period absent adjustment, and the max operator reflects the discrete choice between not adjusting (continuing with $V(x')$) and adjusting (resetting the price gap to $x'' = 0$ at the optimal price, at a cost κ). Define:

$$V_{adj} = V(0) - \kappa, \tag{30}$$

as the value of adjusting. The Bellman equation becomes:

$$V(x) = \pi(x) + \beta \, \mathbb{E}\left[\max\left\{V(x'),\ V_{adj}\right\}\right]. \tag{31}$$

The solution to the Bellman equation (31) takes the form of an (S, s) policy: the firm does not adjust its price as long as the price gap remains within the inaction region (s_L, s_U), and adjusts to the return point $x^* = 0$ whenever the price gap exits this region.

Lemma 1 (Optimal (S, s) Policy) *The optimal pricing policy is characterized by thresholds $s_L < 0 < s_U$ such that:*

$$\textit{Firm adjusts if and only if } x_{i,t} \leq s_L \textit{ or } x_{i,t} \geq s_U.$$

At the adjustment boundaries, the smooth-pasting conditions hold:

$$V'(s_L) = 0, \qquad V'(s_U) = 0,$$

and the value-matching conditions are:

$$V(s_L) = V(0) - \kappa, \qquad V(s_U) = V(0) - \kappa.$$

In the symmetric case where the drift of the optimal price is zero in steady state, the thresholds are symmetric: $s_U = -s_L \equiv \bar{s}$.

5 Pass-Through Analysis

We now analyze how the menu-cost pricing model interacts with the cost structure to determine the pass-through of trade cost shocks to producer prices. We decompose pass-through into its constituent channels and show how the decomposition differs between normal times and crisis periods.

5.1 Shock Structure and the Derivation of the Price Gap Innovation

Each supplier j's shipping cost shock can be decomposed into a route-specific idiosyncratic component and a common aggregate component:

$$\Delta \ln \tau_{j,t} = \underbrace{\nu_{j,t}}_{\substack{\text{idiosyncratic} \\ \text{to route } j}} + \underbrace{u_t}_{\substack{\text{common to} \\ \text{all routes}}}, \tag{32}$$

where $\nu_{j,t} \overset{iid}{\sim} (0, \sigma_\nu^2)$ are independent across suppliers j and time t, and $u_t \sim (0, \sigma_u^2)$ is the aggregate shipping cost shock common to all routes. In normal times, $\sigma_u = 0$ (shipping cost movements are route-specific). During a crisis (pandemic), $\sigma_u > 0$ reflects the synchronized surge in shipping costs across all trade routes, as documented in our empirical results.

From equation (12), taking logs and first-differencing the firm's marginal cost yields:

$$\Delta \ln MC_{i,t} = \underbrace{\Delta \ln C_t^*}_{\substack{\text{common} \\ \text{factor prices}}} - \underbrace{\Delta \ln \Omega_{i,t}}_{\substack{\text{firm-specific} \\ \text{productivity}}} - \underbrace{\phi \, \Delta \ln B_{i,t}}_{\substack{\text{import cost} \\ \text{changes}}}. \tag{33}$$

The third term is where trade cost shocks enter. Given that $B_t = \prod_{j=1}^{N} b_{j,t}^{\gamma_j}$ and $d \ln b_{j,t} / d \ln \tau_{j,t} =$

$-\tilde{a}_{j,t}$, we have

$$\begin{aligned}\Delta \ln MC_{i,t} &= \Delta \ln C_t^* - \Delta \ln \Omega_{i,t} + \phi \sum_{j=1}^{N} \gamma_j \, \tilde{a}_{j,t} \, \Delta \ln \tau_{j,t} \qquad (34)\\ &= \Delta \ln C_t^* - \Delta \ln \Omega_{i,t} + \phi \sum_{j=1}^{N} \gamma_j \, \tilde{a}_{j,t} \, \nu_{j,t} + \phi \left(\sum_{j=1}^{N} \gamma_j \, \tilde{a}_{j,t} \right) u_t \\ &= \Delta \ln C_t^* - \Delta \ln \Omega_{i,t} + \underbrace{\phi \sum_{j=1}^{N} \gamma_j \, \tilde{a}_{j,t} \, \nu_{j,t}}_{\text{idiosyncratic trade cost component}} + \underbrace{\phi \, s_M \cdot u_t}_{\text{aggregate trade cost component}},\end{aligned}$$

where $\tilde{a}_{j,t}$ is the expenditure share of imports in type-j intermediate input and we have used the definition of $s_M \equiv \sum_{j=1}^{N} \gamma_j \, \tilde{a}_{j,t}$ as the aggregate import expenditure share.

Next, we derive how markup changes respond to cost shocks. Recall that the markup $\mu_{i,t} = \sigma_{i,t}/(\sigma_{i,t} - 1)$ depends on the effective demand elasticity $\sigma_{i,t} = \rho(1 - S_{i,t}) + \eta S_{i,t}$, which in turn depends on the firm's market share $S_{i,t} = \xi_{i,s} \left(\frac{p_{i,t}}{P_{s,t}} \right)^{1-\rho}$ and note that the market share responds to the firm's price relative to the sector price index. When the firm resets its price to $p_{i,t}^*$, its market share adjusts, which changes the demand elasticity and hence the desired markup. This creates a fixed-point problem: the optimal price depends on the markup, which depends on the optimal price itself.

Log-linearizing around the symmetric steady state (where $\bar{S}$, $\bar{\sigma}$, and $\bar{\mu}$ denote steady-state values), we can express the markup response to a change in the optimal price:

$$\Delta \ln \mu_{i,t} = -\Gamma \cdot \Delta \ln p_{i,t}^*, \qquad (35)$$

where, note that $(\rho - 1)/(\bar{\sigma} - 1) \approx 1$ for the relevant parameter range,

$$\Gamma \equiv -\frac{d \ln \mu}{d \ln p^*} = \frac{\bar{S}(\rho - \eta)(1 - \rho)}{\bar{\sigma}^2(\bar{\sigma} - 1)} \cdot (-1) \approx \frac{\bar{S}(\rho - \eta)}{\bar{\sigma}(\bar{\sigma} - 1)}. \qquad (36)$$

In deriving the expression for Γ, we have used the findings that $\frac{d \ln \mu}{d \ln \sigma} = \frac{-1}{\sigma(\sigma - 1)}$, $\frac{d \ln \sigma}{d \ln S} = \frac{S(\eta - \rho)}{\sigma}$, $\frac{d \ln S}{d \ln p^*} = 1 - \rho$, and $\frac{d \ln \mu}{d \ln p^*} = \frac{d \ln \mu}{d \ln \sigma} \cdot \frac{d \ln \sigma}{d \ln S} \cdot \frac{d \ln S}{d \ln p^*}$. The negative sign means that when the firm raises its price ($\Delta \ln p^* > 0$), it loses market share, faces more elastic demand, and its desired markup *falls*.

Lastly, we turn to the derviation of the law of motion for the firm's optimal price $p_{i,t}^*$ and show how it decomposes into idiosyncratic and aggregate components. This derivation connects the static cost structure to the dynamic pricing problem. First-differencing between

periods $t-1$ and t in equation (22) and substituting equation (35) yields:

$$\begin{aligned}\Delta \ln p^*_{i,t} &= \Delta \ln \mu_{i,t} + \Delta \ln MC_{i,t} = \frac{1}{1+\Gamma} \Delta \ln MC_{i,t} \\ &= \frac{1}{1+\Gamma} \left[\Delta \ln C^*_t - \Delta \ln \Omega_{i,t} + \phi \sum_{j=1}^{N} \gamma_j \, \tilde{a}_{j,t} \, \nu_{j,t} + \phi \, s_M \, u_t \right] \\ &\equiv \varepsilon^{idio}_{i,t} + \varepsilon^{agg}_t,\end{aligned}$$

where we define the idiosyncratic and aggregate components:

$$\varepsilon^{idio}_{i,t} \equiv \frac{1}{1+\Gamma} \left[-\Delta \ln \Omega_{i,t} + \phi \sum_{j=1}^{N} \gamma_j \, \tilde{a}_{j,t} \, \nu_{j,t} \right], \tag{37}$$

$$\varepsilon^{agg}_t \equiv \frac{1}{1+\Gamma} \left[\Delta \ln C^*_t + \phi \, s_M \, u_t \right]. \tag{38}$$

Note $\varepsilon^{idio}_{i,t}$ is firm-specific, driven by idiosyncratic productivity shocks ($\Delta \ln \Omega_{i,t}$) and the diversified average of route-specific shipping cost shocks ($\bar{\nu}_{i,t}$). Its variance from the trade cost channel is proportional to $(\phi s_M)^2 \sigma^2_\nu / N$. On the other hand, ε^{agg}_t is common to all firms, driven by aggregate factor prices ($\Delta \ln C^*_t$) and the correlated shipping cost shock (u_t). Its variance from the trade cost channel is proportional to $(\phi s_M)^2 \sigma^2_u$ — a factor of N larger than the idiosyncratic component for shocks of equal size.

In normal times, shipping cost shocks are route-specific ($\sigma_u = 0$, $\sigma_\nu > 0$). The aggregate component ε^{agg}_t contains only $\Delta \ln C^*_t$, which is small in partial equilibrium. The firm-level optimal price is driven primarily by $\varepsilon^{idio}_{i,t}$, which has a small trade cost variance due to the $1/N$ diversification effect.

During the pandemic ($\sigma_u > 0$), correlated shipping cost shocks activate the aggregate component, and the effective variance of shipping-driven optimal price changes is N times larger. Combined with the nonlinearity of the menu cost friction, this generates the amplified pass-through observed in the data.

The ratio of crisis-to-normal effective shipping cost variance for a firm with N suppliers is:

$$\frac{\text{Var}(\varepsilon^{agg} \mid \text{shipping})}{\text{Var}(\varepsilon^{idio} \mid \text{shipping})} = \frac{(\phi s_M)^2 \sigma^2_u}{(\phi s_M)^2 \sigma^2_\nu / N} = N \cdot \frac{\sigma^2_u}{\sigma^2_\nu}.$$

For $N = 5$ (our calibrated value based on the LFTTD data showing an average of 5.3 source countries) and shocks of comparable magnitude ($\sigma_u \approx \sigma_\nu$), the effective variance of shipping cost shocks to marginal cost is 5 times larger during a crisis with perfectly correlated shocks than during normal times with independent shocks. This is the fundamental mechanism generating amplified pass-through.

5.2 Pass-Through Decomposition

Consider a shipping cost shock $d\ln\tau$ that is either idiosyncratic (affecting one supplier) or aggregate (affecting all suppliers). The total pass-through of this shock to the firm's output price can be decomposed into three multiplicative channels:

$$\frac{d\ln p_i}{d\ln\tau} = \underbrace{\frac{d\ln MC_i}{d\ln\tau}}_{\text{(I) Direct cost}} \times \underbrace{\frac{d\ln p_i^*}{d\ln MC_i}}_{\text{(II) Markup absorption}} \times \underbrace{\frac{d\ln p_i}{d\ln p_i^*}}_{\text{(III) Menu cost friction}} . \tag{39}$$

Each component operates through a distinct economic mechanism and contributes differently to pass-through attenuation depending on whether the shock is idiosyncratic or aggregate. We now discuss each in turn.

The first component captures the direct cost channel and measures how much a shipping cost shock raises the firm's marginal cost. From Proposition 2, this elasticity depends critically on whether the shock is idiosyncratic or aggregate:

$$\frac{d\ln MC_i}{d\ln\tau} = \begin{cases} \alpha_d\,\omega_j & \text{idiosyncratic shock to supplier } j, \\ \alpha_d & \text{common shock to all suppliers,} \end{cases} \tag{40}$$

where $\alpha_d = \phi s_M$ is the overall cost share of imported intermediates in the firm's production—the product of the intermediate input share ϕ and the import expenditure share s_M—and $\omega_j \approx 1/N$ is the share of supplier j in total import expenditure, which equals $1/N$ for symmetric suppliers.

The distinction between idiosyncratic and aggregate shocks is the first source of the normal-times versus crisis asymmetry. When a shipping cost shock hits only one of the firm's N suppliers, its impact on marginal cost is diluted by a factor of $\omega_j \approx 1/N$, because the affected route accounts for only a fraction of the firm's total import bill. This is the *diversification effect*: multi-sourcing insulates the firm from route-specific disturbances by spreading exposure across N independent supply chains. In contrast, when a common shock hits all suppliers simultaneously—as occurred during the pandemic—the diversification benefit vanishes entirely. The full import cost share α_d is exposed, and the marginal cost elasticity is N times larger than in the idiosyncratic case. This N-fold amplification of the direct cost channel is the first mechanism through which correlated shocks generate disproportionately larger pass-through.

The second component captures the extent to which a change in marginal cost is reflected in the firm's frictionless optimal price p_i^* (i.e., markup absorption). Under the nested CES demand structure, the optimal price satisfies $p_i^* = \mu_i \cdot MC_i$, where μ_i is the desired markup.

The elasticity of the optimal price with respect to marginal cost is therefore:

$$\frac{d \ln p_i^*}{d \ln MC_i} = \frac{d \ln \mu_i}{d \ln MC_i} + 1 = 1 - \Gamma, \tag{41}$$

where

$$\Gamma = \frac{\bar{S}(\rho - \eta)}{\bar{\sigma}(\bar{\sigma} - 1)} \tag{42}$$

is the markup elasticity with respect to cost changes.

This component is less than one whenever $\Gamma > 0$, reflecting the strategic complementarity inherent in the nested demand structure. When a cost increase raises the firm's optimal price, the firm loses market share relative to its within-sector competitors. This loss of market share shifts the firm's effective demand elasticity $\sigma_i = \rho(1 - S_i) + \eta S_i$ toward the higher within-sector elasticity ρ, making demand more elastic and compressing the desired markup. The firm therefore finds it optimal to absorb part of the cost increase through a lower markup rather than passing the entire increase on to its price.

The magnitude of this absorption depends on the gap between the within-sector and across-sector elasticities $(\rho - \eta)$ and on the firm's initial market share $\bar{S}$. Firms with larger market shares face a steeper trade-off between pricing and market share, and consequently exhibit greater markup absorption. For the median firm in our calibration, $\Gamma \approx 0.12$, implying that approximately 12% of any marginal cost increase is absorbed through markup compression, with the remaining 88% passed through to the optimal price. Unlike the direct cost channel, the markup absorption rate $(1 - \Gamma)$ is the same regardless of whether the underlying shock is idiosyncratic or aggregate—it depends only on the firm's competitive position, not on the nature of the cost disturbance. Markup absorption therefore attenuates pass-through uniformly across both regimes.

The third component introduces the key nonlinearity due to menu costs that distinguishes the model from standard flexible-price frameworks. It measures how much of the change in the frictionless optimal price p_i^* is actually reflected in the firm's posted price. Define the menu cost friction as:

$$\mathcal{F}(\Delta \ln p^*) \equiv \frac{d \ln \bar{p}}{d \ln p^*}, \tag{43}$$

where $\bar{p}$ is the average posted price across firms. This friction depends crucially on the *size* of the optimal price change $\Delta \ln p^*$ relative to the width of the inaction region $\bar{s}$.

The economic intuition is straightforward. In a menu cost model, firms tolerate small deviations between their posted price and their optimal price rather than paying the lump-sum adjustment cost. Only when the price gap—the log deviation of the posted price from its optimal level—crosses the boundary of the (S, s) inaction region does the firm find it

worthwhile to reset its price. The aggregate price response therefore depends on what fraction of firms are pushed outside their inaction regions by the shock, which in turn depends on the shock's magnitude.

When the optimal price change is small relative to $\bar{s}$, only the fraction Λ of firms that would have adjusted in steady state (plus a small selection margin of newly triggered firms) actually change their prices. Most firms absorb the shock within their inaction bands, and the aggregate price barely moves. When the optimal price change is large—exceeding $2\bar{s}$—every firm is pushed outside its inaction region and adjusts fully, so $\mathcal{F} = 1$. Between these extremes, $\mathcal{F}$ is strictly increasing in the shock size: larger shocks trigger progressively more firms to adjust, generating a convex aggregate price response. We formalize these properties in the following proposition.

Proposition 3 (Menu Cost Pass-Through) *Let $\delta \equiv |\Delta \ln p^*|$ denote the magnitude of the common shift in the optimal price. The menu cost friction $\mathcal{F}(\delta) \equiv \Delta\bar{p}/\delta$ depends on the size of the shock relative to the inaction band:*

1. *For small shocks ($\delta \ll \bar{s}$), $\mathcal{F}(\delta) \approx \Lambda + \Lambda'$, where Λ is the steady-state frequency of adjustment and $\Lambda' = f(-\bar{s}) \cdot \bar{s} > 0$ is the selection effect—the additional pass-through from firms newly pushed outside their inaction regions. Since both Λ and Λ' are small (of order 0.10 in our calibration), $\mathcal{F}$ is small.*

2. *For large shocks ($\delta \geq 2\bar{s}$), $\mathcal{F}(\delta) = 1$: all firms are pushed outside their inaction regions and adjust their prices fully.*

3. *$\mathcal{F}(\delta)$ is continuous and strictly increasing in δ on $[0, 2\bar{s}]$.*

5.3 Quantifying Pass-Through

We calibrate the model at a monthly frequency. The parameters are divided into four groups: preference and technology, trade and sourcing, menu-cost pricing, and shock processes. Table 4 summarizes the calibration. Details about the calibraiton are in Appendix B.

Normal-Times Pass-Through. In normal times, a 10% shipping cost increase to one supplier generates a change in the firm's optimal price of:

$$\begin{aligned}\Delta \ln p_i^* = \alpha_d \, \omega_j \, (1 - \Gamma) \, \Delta \ln \tau_j &\approx \frac{\phi s_M}{N} (1 - \Gamma) \, \Delta \ln \tau_j \\ &= \frac{0.50 \times 0.20}{5} \times 0.88 \times 0.10 = 0.0018,\end{aligned}$$

TABLE 4: Model Calibration

Parameter	Symbol	Value	Target/Source
Panel A: Preferences and Technology			
Discount factor	β	0.9967	4% annual real rate
Intermediate input share	ϕ	0.50	U.S. manufacturing I-O tables
Across-sector elasticity	η	3.0	Atkeson-Burstein (2008)
Within-sector elasticity	ρ	8.0	≈21% avg. markup
Average demand elasticity	$\bar{\sigma}$	5.75	$\rho(1-\bar{S})+\eta\bar{S}$
Steady-state market share	$\bar{S}$	0.30	PPI microdata
Panel B: Trade and Sourcing			
Number of source countries	N	5	LFTTD: 5.3 avg. source countries
Armington elasticity	ζ	2.0	Literature estimates
Cobb-Douglas weights	γ_j	0.20	Symmetric $(1/N)$
Steady-state import share	s_M	0.20	Median, LFTTD-PPI sample
SS ad valorem shipping cost	$\bar{\tau}-1$	4.82%	Table 1, Panel A
Crisis shipping cost	$\tau^{crisis}-1$	≈8%	Figure 1
Panel C: Menu Cost Pricing			
Menu cost	κ	0.008	≈10% monthly adj. freq.
Inaction band half-width	$\bar{s}$	0.062	Implied by κ
Panel D: Shock Processes—Normal Times			
Gaussian std. dev.	σ_G	0.025	PPI price change distribution
Poisson arrival rate	λ_p	0.02	Kurtosis of price changes
Poisson jump std. dev.	σ_p	0.05	Kurtosis of price changes
Aggregate shock std. dev.	σ_{agg}	0	No aggregate shocks
Panel E: Shock Processes—Crisis (Pandemic)			
Poisson arrival rate	λ_p^{crisis}	0.08	Increased volatility
Poisson jump std. dev.	σ_p^{crisis}	0.08	Larger shocks
Aggregate shock std. dev.	σ_{agg}^{crisis}	0.04	IV estimates, Table 1

where we have used our calibrated parameter values ($\phi = 0.50$, $s_M = 0.20$, $N = 5$, $\Gamma = 0.12$). This is well within the inaction region (which has a half-width of approximately $\bar{s} \approx 0.06$ or 6% in our calibration), so the menu cost friction $\mathcal{F}$ is approximately equal to $\Lambda \approx 0.10$.

The total pass-through in normal times is therefore:

$$\gamma_{normal} \approx \frac{\phi s_M}{N} \times (1-\Gamma) \times \Lambda \approx \frac{0.10}{5} \times 0.88 \times 0.10 = 0.0018, \tag{44}$$

which is economically negligible and statistically indistinguishable from zero in a regression

with the typical sample sizes in our data. This matches the empirical finding of near-zero pass-through during normal times.

Crisis Pass-Through. During the pandemic, shipping costs rose simultaneously across all trade routes and source countries. The change in the firm's optimal price is:

$$\begin{aligned}\Delta \ln p_i^* &= \alpha_d \left(1-\Gamma\right) \Delta \ln \tau = \phi s_M \left(1-\Gamma\right) \Delta \ln \tau \\ &= 0.10 \times 0.88 \times 0.56 \approx 0.049,\end{aligned}$$

where the key difference is that the shock hits *all* N suppliers simultaneously, eliminating the $1/N$ diversification effect. That is, a 75% increase in shipping costs generates a price increase of about 5%. This is close to the inaction band ($\bar{s} \approx 0.06$), so $\mathcal{F}$ is in the intermediate regime—substantially above Λ but below 1.

Total crisis pass-through. For the pandemic episode, combining the cost share with the markup absorption and menu cost friction, the total pass-through during the pandemic is:

$$\gamma_{crisis} \approx \alpha_d \left(1-\Gamma\right) \mathcal{F}(\Delta \ln p^*_{crisis}) \approx 0.10 \times 0.88 \times 0.41 \approx 0.036, \tag{45}$$

where $\mathcal{F} \approx 0.41$ reflects that the correlated shock pushes a substantial fraction of firms outside their inaction regions, but not all: the shock ($\delta \approx 0.049$) is approximately $0.8\bar{s}$, placing it in the intermediate regime of the menu-cost friction function (Proposition 3). The pass-through coefficient of approximately 0.036 at the 12-month horizon is consistent with our IV estimates.

Simulation-based Quantitative Analysis. We solve the model numerically using value function iteration on a fine grid over the price gap x (see Appendix C for algorithmic details). We then simulate the model under both normal-time and crisis shock processes to generate model-implied pass-through coefficients and impulse response functions.

Table 5 compares the model-implied pass-through coefficients to the empirical estimates from the companion paper. We report both the normal-times baseline (γ) and the crisis interaction (γ_{crisis}) at three horizons.

The model successfully matches the key patterns in the data. First, it reproduces near-zero normal-times pass-through. The model generates baseline pass-through coefficients (γ) that are economically small and statistically insignificant at all horizons, consistent with the data. The mechanism is the triple attenuation: idiosyncratic shocks are diversified across N

TABLE 5: SPass-Through Coefficients: Model vs. Data

	Data (IV, Firm FE)			Model		
Horizon	3 mo	6 mo	12 mo	3 mo	6 mo	12 mo
γ (normal)	0.007	0.013	0.022	0.002	0.005	0.009
	(0.012)	(0.020)	(0.036)			
γ_{crisis} (pandemic)	0.011	0.021	0.036	0.011	0.022	0.036
	(0.005)	(0.008)	(0.014)			

NOTE:Data columns report IV estimates from Table 1, Panel B, columns (4)–(6) of the companion paper, with standard errors in parentheses. Model columns report simulated pass-through coefficients from the calibrated menu-cost model. The model coefficients are generated by simulating 10,000 firms for 240 months, imposing a shipping cost shock of $\Delta \ln \tau = 0.56$ (corresponding to the observed pandemic-era increase), and running the same regression specification as in the empirical analysis.

suppliers, markup absorption further reduces the optimal price response, and the resulting small change in the optimal price falls within firms' inaction regions.

Second, it generates amplified crisis pass-through. The model generates crisis interaction coefficients (γ_{crisis}) of 0.011 at 3 months, 0.022 at 6 months, and 0.036 at 12 months, closely matching the empirical estimates of 0.011, 0.021, and 0.036. The amplification comes from two sources: (i) correlated shocks eliminate diversification, multiplying the effective cost impact by N; and (ii) the larger shock pushes a substantial fraction of firms outside their inaction regions, increasing the menu-cost friction $\mathcal{F}$ from approximately 0.10 to approximately 0.41.

Lastly, the model generates pass-through that is increasing in the horizon, with the ratio of 12-month to 3-month pass-through approximately 3.2 in the model versus 3.3 in the data. This pattern arises because menu costs slow the adjustment process: in each period after the shock, an additional fraction of firms cross their adjustment thresholds and reset their prices closer to the new optimal level.

6 Conclusion

This paper provides new evidence on how trade costs pass through to U.S. producer prices, exploiting a novel panel dataset that links confidential firm-level trade transactions to confidential firm-level output prices. Our main finding is that the pass-through of shipping costs to producer prices is highly state-dependent: negligible during normal times, but substantially amplified during the COVID-19 pandemic.

Several mechanisms underlie this asymmetry. During normal times, firms can insulate their output prices from shipping cost fluctuations by substituting across suppliers, adjusting markups, and renegotiating contracts. During the pandemic, these adjustment margins were

severely curtailed by the broad-based and highly correlated nature of the shipping cost surge, forcing firms to pass a larger share of cost increases through to output prices. Our theoretical framework formalizes this intuition by embedding multi-country sourcing into a menu-cost pricing model with variable markups. The model decomposes pass-through into three channels—a direct cost channel attenuated by supplier diversification, markup absorption due to strategic complementarity, and a menu-cost friction that is increasing in shock size—and shows that the state-dependent pass-through arises naturally from the interaction of supplier diversification and nonlinear price adjustment. In normal times, idiosyncratic route-level shocks are diversified across suppliers and produce optimal price changes that fall within firms' inaction regions, yielding negligible pass-through. During the pandemic, correlated shocks across all trade routes eliminate diversification and push firms outside their inaction regions, generating the amplified pass-through observed in the data. The calibrated model quantitatively matches the empirical pass-through at all horizons.

Our findings have important implications for understanding the inflationary consequences of trade shocks. The results suggest that the degree to which trade cost increases translate into domestic inflation depends critically on the nature of the shock: idiosyncratic, targeted shocks (such as tariffs on specific products from specific countries) are likely to be largely absorbed by firms, whereas broad-based, aggregate shocks (such as global supply chain disruptions) can have significant inflationary effects. This distinction is relevant for policymakers assessing the potential price impact of both trade policy changes and supply chain vulnerabilities.

The state-dependent nature of pass-through also has implications for monetary policy. Our results suggest that central banks should pay particular attention to the breadth and correlation structure of trade cost shocks when assessing their inflationary potential. A broad-based increase in trade costs—whether due to supply chain disruptions, geopolitical events, or sweeping tariff policies—could generate more significant and persistent inflationary pressure.

Looking forward, our framework suggests that recent and prospective tariff increases on a broad set of U.S. trading partners could generate significantly larger pass-through than the more targeted tariffs of 2018–2019, to the extent that they limit firms' ability to substitute across source countries. Quantifying the inflationary effects of such broad-based tariff regimes, and understanding how firms adjust their sourcing strategies in response, are important avenues for future research as new data become available.

References

Amiti, M., S. Heise, and D. Weinstein (2023): "The Effect of the Strong Dollar on U.S. Import Prices and Inflation," *AEA Papers and Proceedings*, 113, 541–545.

Amiti, M., O. Itskhoki, and J. Konings (2019): "International Shocks, Variable Markups, and Domestic Prices," *Review of Economic Studies*, 86, 2356–2402.

Amiti, M., S. J. Redding, and D. Weinstein (2019): "The Impact of the 2018 Tariffs on Prices and Welfare," *Journal of Economic Perspectives*, 33, 187–210.

Antràs, P. and A. Costinot (2011): "Intermediated Trade," *Quarterly Journal of Economics*, 126, 1319–1374.

Atkeson, A. and A. Burstein (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98, 1998–2031.

Autor, D. H., D. Dorn, and G. H. Hanson (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–2168.

Bai, X., J. Fernández-Villaverde, Y. Li, and F. Zanetti (2025): "The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory," NBER Working Paper No. 32098.

Baldacchino, G., R. Duval, and others (2023): "Supply Chain Disruptions and Inventory Dynamics," *IMF Working Paper* WP/23/45.

Berman, N., P. Martin, and T. Mayer (2012): "How Do Different Exporters React to Exchange Rate Changes?," *Quarterly Journal of Economics*, 127, 437–492.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2018): "Global Firms," *Journal of Economic Literature*, 56, 565–619.

Bonadio, B., Z. Huo, A. A. Levchenko, and N. Pandalai-Nayar (2021): "Global Supply Chains in the Pandemic," *Journal of International Economics*, 133, 103534.

Campa, J. M. and L. S. Goldberg (2005): "Exchange Rate Pass-Through into Import Prices," *Review of Economics and Statistics*, 87, 679–690.

Cavallo, A., G. Gopinath, B. Neiman, and J. Tang (2021): "Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy," *American Economic Review: Insights*, 3, 19–34.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016): "Prices, Markups, and Trade Reform," *Econometrica*, 84, 445–510.

Di Giovanni, J., Ş. Karakaş, A. Ozhan, and A. Sala (2022): "Global Supply Chain Pressures, International Trade, and Inflation," *NBER Working Paper* No. 30240.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020): "The Return to Protectionism," *Quarterly Journal of Economics*, 135, 1–55.

Flaaen, A., A. Hortaçsu, and F. Tintelnot (2020): "The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines," *American Economic Review*, 110, 2103–2127.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek (2017): "Inflation Dynamics During the Financial Crisis," *American Economic Review*, 107, 785–823.

Golosov, M. and R. E. Lucas, Jr. (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–199.

Goldberg, P. K. and M. M. Knetter (1997): "Goods Prices and Exchange Rates: What Have We Learned?," *Journal of Economic Literature*, 35, 1243–1272.

Gopinath, G., O. Itskhoki, and R. Rigobon (2010): "Currency Choice and Exchange Rate Pass-Through," *American Economic Review*, 100, 304–336.

Handley, K. and N. Limão (2017): "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, 107, 2731–2783.

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.

Kamal, F. and W. Ouyang (2020): "Identifying U.S. Merchandise Traders: Integrating Customs Transactions with Business Administrative Data," *CES Working Paper* CES-20-28.

Kimball, M. S. (1995): "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, 27, 1241–1277.

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.

LeBlanc, J. and R. Bhatt (2024): "Global Supply Chain Disruptions and Domestic Inflation," *International Finance Discussion Papers*, No. 1395.

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

Pierce, J. R. and P. K. Schott (2016): "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, 106, 1632–1662.

Shapiro, J. S. (2016): "Trade Costs, CO2, and the Environment," *American Economic Journal: Economic Policy*, 8, 220–254.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020): "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 110, 2586–2624.

Borusyak, K., P. Hull, and X. Jaravel (2022): "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 89, 181–213.

A Proofs

Proof of Proposition 1.

We derive the total variable cost function in the following steps.

Step 1: Optimal input mix for given composite intermediate. The firm's cost minimization problem is:

$$\min_{X_{i,t},L_{i,t}} \tilde{V}_t X_{i,t} + W_t L_{i,t} \quad \text{s.t.} \quad Y_{i,t} = \Omega_{i,t} X_{i,t}^{\phi} L_{i,t}^{1-\phi}. \tag{46}$$

The first-order conditions yield:

$$\tilde{V}_t = \lambda\phi \frac{Y_{i,t}}{X_{i,t}}, \tag{47}$$

$$W_t = \lambda(1-\phi) \frac{Y_{i,t}}{L_{i,t}}, \tag{48}$$

where λ is the Lagrange multiplier (marginal cost). Substituting the factor demands back into the production function, we obtain the marginal cost:

$$\lambda = \frac{1}{\Omega_{i,t}} \left(\frac{\tilde{V}_t}{\phi}\right)^{\phi} \left(\frac{W_t}{1-\phi}\right)^{1-\phi}. \tag{49}$$

The total variable cost is $TVC = \lambda \cdot Y_{i,t}$.

Step 2: CES unit cost for each input type. For type-j intermediates, the CES unit cost index from (6) is:

$$\tilde{V}_{j,t} = \frac{V_{j,t}^*}{b_{j,t}}, \tag{50}$$

since all N countries are active. The cost-reduction factor from importing type-j intermediates is:

$$b_{j,t} = \left[1 + a_j \left(\frac{P_{j,t}^M}{V_{j,t}^*}\right)^{-\zeta}\right]^{1/\zeta}.$$

To verify (50), note that the CES unit cost of the home-foreign composite (6) is:

$$\tilde{V}_{j,t} = \left[(V_{j,t}^*)^{-\zeta} + a_j (P_{j,t}^M)^{-\zeta}\right]^{-1/\zeta} = V_{j,t}^* \left[1 + a_j \left(\frac{P_{j,t}^M}{V_{j,t}^*}\right)^{-\zeta}\right]^{-1/\zeta} = \frac{V_{j,t}^*}{b_{j,t}}. \tag{51}$$

Note that $b_{j,t} > 1$ whenever the imported variety is available and the quality-adjusted import price is finite, reflecting the gains from import variety.

Step 3: Cobb-Douglas aggregation and total variable cost. The effective price index for the composite intermediate is the Cobb-Douglas aggregate:

$$\tilde{V}_t = \prod_{j=1}^{N} \left(\frac{\tilde{V}_{j,t}}{\gamma_j} \right)^{\gamma_j} = \frac{V_t^*}{B_{N,t}}, \tag{52}$$

where $V_t^* \equiv \prod_{j=1}^{N} (V_{j,t}^*/\gamma_j)^{\gamma_j}$ is the Cobb-Douglas price index over domestic input types, and the cumulative cost-reduction factor from importing from all N countries is:

$$B_{N,t} = \prod_{j=1}^{N} b_{j,t}^{\gamma_j}.$$

Substituting (52) into (49), the total variable cost function is:

$$TVC_t(Y_{i,t}) = \frac{C_t^*}{B_{N,t}^{\phi}} \cdot \frac{Y_{i,t}}{\Omega_{i,t}},$$

where the unit cost index for a non-importing firm is:

$$C_t^* = \left(\frac{V_t^*}{\phi} \right)^{\phi} \left(\frac{W_t}{1-\phi} \right)^{1-\phi}.$$

■

Derivation of Flow Profit in Equation (28). We first express $\Pi_{i,t}$ as a function of the price gap $x_{i,t} = \ln p_{i,t} - \ln p_{i,t}^*$, where $p_{i,t}^* = \mu_{i,t} \cdot MC_{i,t}$ is the optimal price from (22). Define the relative price $\hat{p}_{i,t} \equiv p_{i,t}/p_{i,t}^*$, so $\ln \hat{p}_{i,t} = x_{i,t}$, i.e., $\hat{p}_{i,t} = e^{x_{i,t}}$.

Revenue is $p_{i,t} \cdot q_{i,t}$. From the demand function, holding $P_{s,t}$ and D_t fixed:

$$q_{i,t} \propto p_{i,t}^{-\sigma}, \tag{53}$$

where $\sigma = \bar{\sigma}$ is the effective demand elasticity evaluated at the steady-state market share (we use the local approximation around steady state, holding σ constant). Therefore revenue is proportional to $p_{i,t}^{1-\sigma}$. At the optimal price p^*, revenue is proportional to $(p^*)^{1-\sigma}$. The ratio is:

$$\frac{p_{i,t} \cdot q(p_{i,t})}{p_{i,t}^* \cdot q(p_{i,t}^*)} = \left(\frac{p_{i,t}}{p_{i,t}^*} \right)^{1-\sigma} = e^{(1-\sigma)x}. \tag{54}$$

For total cost, we have $TC_{i,t} = MC_{i,t} \cdot q_{i,t}$ (since marginal cost is constant in output under the Cobb-Douglas technology). The quantity produced is $q_{i,t} \propto p_{i,t}^{-\sigma}$, so:

$$\frac{MC_{i,t} \cdot q(p_{i,t})}{MC_{i,t} \cdot q(p_{i,t}^*)} = \left(\frac{p_{i,t}}{p_{i,t}^*} \right)^{-\sigma} = e^{-\sigma x}. \tag{55}$$

Lastly, we normalize flow profit by steady-state revenue $p^* \cdot q(p^*)$. The normalized profit

at price gap x is:

$$\begin{aligned}\pi(x) &= \frac{p_{i,t} \cdot q(p_{i,t}) - MC_{i,t} \cdot q(p_{i,t})}{p^* \cdot q(p^*)} \\ &= e^{(1-\sigma)x} - \frac{MC_{i,t}}{p^*} \cdot e^{-\sigma x}. \end{aligned} \tag{56}$$

At the optimal price, $p^* = \mu \cdot MC$, so $MC/p^* = 1/\mu = (\sigma - 1)/\sigma$. Substituting and writing $\theta \equiv \bar{\sigma}$ yields equation (28).

Properties of the flow profit function. We verify:

$$\begin{aligned}\pi(0) &= 1 - \frac{\theta - 1}{\theta} = \frac{1}{\theta}, \\ \pi'(x) &= (1-\theta)e^{(1-\theta)x} + (\theta - 1)e^{-\theta x}, \\ \pi'(0) &= (1-\theta) + (\theta - 1) = 0, \\ \pi''(x) &= (1-\theta)^2 e^{(1-\theta)x} - \theta(\theta - 1)e^{-\theta x}, \\ \pi''(0) &= 1 - \theta. \end{aligned}$$

Since $\theta = \bar{\sigma} > 1$, we have $\pi''(0) = 1 - \theta < 0$, confirming the profit function is concave at $x = 0$ ■

Proof of Proposition 3. Let $f(x)$ be the stationary density of price gaps on $[-\bar{s}, \bar{s}]$ and consider a positive common shock $\delta > 0$ that raises the optimal price p^*. Since $x = \ln p - \ln p^*$, the price gap of every firm decreases: $x_i' = x_i - \delta$. This pushes all firms toward the *lower* boundary $-\bar{s}$.

Part (i). For small δ, only firms with pre-shock gaps $x_i \in [-\bar{s}, -\bar{s} + \delta]$ are newly pushed below $-\bar{s}$. Their mass is $\approx f(-\bar{s}) \cdot \delta$. Each resets to $x = 0$, generating a price increase of $\approx \bar{s}$. This contributes $f(-\bar{s}) \cdot \bar{s} \cdot \delta$ to the aggregate price change. Adding the steady-state adjusters' contribution $\Lambda \cdot \delta$ and dividing by δ: $\mathcal{F} \approx \Lambda + f(-\bar{s}) \cdot \bar{s}$.

Part (ii). For $\delta \geq 2\bar{s}$: $x_i' = x_i - \delta \leq \bar{s} - 2\bar{s} = -\bar{s}$ for all i, so every firm adjusts. Each firm's price change is $\delta - x_i$, and $\mathbb{E}[\delta - x_i] = \delta$ since $\mathbb{E}[x_i] = 0$ by symmetry. Hence $\mathcal{F} = 1$.

Part (iii). The aggregate price response $\Delta\bar{p}(\delta) = \int_{-\bar{s}}^{-\bar{s}+\delta} (\delta - x) f(x)\, dx + \Lambda\delta$ is convex in δ (since differentiation via Leibniz's rule shows the marginal pass-through is increasing). For a convex function with $\Delta\bar{p}(0) = 0$, the ratio $\mathcal{F}(\delta) = \Delta\bar{p}/\delta$ is increasing.

The detailed derivation with all intermediate steps is in Appendix A. ■

B Calibration

The discount factor $\beta = 0.9967$ corresponds to a 4% annual real interest rate. The intermediate input share $\phi = 0.50$ is taken from the U.S. input-output tables for the manufacturing sector, which show that intermediate inputs account for approximately half of gross output. The across-sector elasticity of substitution $\eta = 3$ and the within-sector elasticity $\rho = 8$ are chosen to generate an average markup of approximately 21%, consistent

with estimates from the PPI data in Gilchrist, Schoenle, Sim, and Zakrajšek (2017). The average effective demand elasticity in steady state is $\bar{\sigma} = \bar{\rho}(1 - \bar{S}) + \eta\bar{S} \approx 5.75$, yielding $\bar{\mu} = 5.75/4.75 \approx 1.21$. The steady-state market share $\bar{S} = 0.30$ implies a markup elasticity $\Gamma = 0.30 \times (8 - 3)/(5.75 \times 4.75) \approx 0.055$.

The number of source countries $N = 5$ is set to match the LFTTD data, which shows an average of 5.3 source countries per importing firm (Table 1, Panel A, of the companion paper). The Armington elasticity $\zeta = 2.0$ is within the range of estimates in the literature, spanning from low values (around 1) for differentiated goods to high values (around 4) for commodities. The Cobb-Douglas weights $\gamma_j = 1/N = 0.20$ assume symmetric input types. The efficiency parameter a_j is calibrated so that the steady-state import share $s_M = \sum_j \gamma_j \tilde{a}_j = 0.40$, matching the average import intensity of firms in our LFTTD-PPI matched sample. The steady-state shipping cost $\bar{\tau}$ is set to generate a 4.82% ad valorem rate, matching the pre-pandemic mean in Table 1, Panel A.

The menu cost $\kappa = 0.008$ is calibrated to match a monthly price adjustment frequency of approximately 10%, consistent with Nakamura and Steinsson (2008). This implies that 90% of firms do not change their nominal price in any given month, and the average duration of a price spell is approximately 10 months. The resulting inaction band has a half-width of approximately $\bar{s} = 0.062$ (6.2%).

The Gaussian component of idiosyncratic shocks has standard deviation $\sigma_G = 0.025$, calibrated to match the standard deviation of monthly log price changes conditional on adjustment. The Poisson arrival rate in normal times is $\lambda_p = 0.02$ (one large shock per 50 months on average) and increases to $\lambda_p^{crisis} = 0.08$ during the pandemic. The Poisson jump size has standard deviation $\sigma_p = 0.05$ in normal times and $\sigma_p^{crisis} = 0.08$ during the pandemic. The aggregate shock standard deviation is $\sigma_{agg} = 0$ in normal times and $\sigma_{agg} = 0.04$ during the pandemic.

C Value Function Iteration Algorithm

We solve the firm's menu-cost pricing problem (31) numerically using value function iteration on a discrete grid. The algorithm proceeds as follows.

C.1 Grid Construction

1. Define a grid of $M = 1000$ equally spaced points over the price gap: $\{x_1, x_2, \ldots, x_M\}$ with $x_1 = -0.20$ and $x_M = 0.20$.

2. Define a grid of shock realizations. For the Gaussian component, use Gauss-Hermite quadrature with $K_G = 15$ nodes $\{z_1^G, \ldots, z_{K_G}^G\}$ and weights $\{w_1^G, \ldots, w_{K_G}^G\}$. For the Poisson-Gaussian component, condition on the number of jumps $n = 0, 1, 2, 3$ (truncate at 3, which accounts for 99.99% of the probability under $\lambda_p \leq 0.08$) and use $K_P = 10$ nodes for the jump-size distribution.

C.2 Iteration Procedure

1. **Initialize.** Set $V^{(0)}(x_m) = \pi(x_m)/(1 - \beta)$ for all grid points m.

2. **Compute continuation value.** For each grid point x_m, compute the expected continuation value:

$$EV^{(n)}(x_m) = \sum_{k_G=1}^{K_G} \sum_{n_p=0}^{3} \Pr(n_p) \sum_{k_P=1}^{K_P} w_{k_G}^G \cdot w_{k_P}^P \cdot \max\left\{V^{(n)}(x'_m),\ V^{(n)}(0) - \kappa\right\}, \quad (57)$$

where $x'_m = x_m - z_{k_G}^G - \sum_{i=1}^{n_p} z_{k_P,i}^P$ and we use linear interpolation for off-grid values of $V^{(n)}(x'_m)$.

3. **Update value function.**

$$V^{(n+1)}(x_m) = \pi(x_m) + \beta\, EV^{(n)}(x_m). \quad (58)$$

4. **Check convergence.** Stop if $\max_m |V^{(n+1)}(x_m) - V^{(n)}(x_m)| < 10^{-10}$.

5. **Extract policy function.** The adjustment thresholds s_L and s_U are determined by:

$$s_L = \max\{x_m : V^*(x_m) \leq V^*(0) - \kappa\}, \qquad s_U = \min\{x_m : V^*(x_m) \leq V^*(0) - \kappa\}, \quad (59)$$

where V^* is the converged value function.

C.3 Simulation

After solving for the value function and policy thresholds, we simulate $I = 10{,}000$ firms for $T = 240$ months:

1. Initialize all firms at $x_{i,0} = 0$.
2. For each period $t = 1, \ldots, T$:
   (a) Draw idiosyncratic shocks $\varepsilon_{i,t}^G \sim N(0, \sigma_G^2)$ and Poisson jumps.
   (b) If in crisis regime, draw common shock $\varepsilon_t^{agg} \sim N(0, \sigma_{agg}^2)$.
   (c) Update price gaps: $x_{i,t} = x_{i,t-1} - \varepsilon_{i,t}^{idio} - \varepsilon_t^{agg}$.
   (d) If $x_{i,t} \leq s_L$ or $x_{i,t} \geq s_U$, the firm adjusts: record $\Delta p_{i,t} = -x_{i,t}$ and set $x_{i,t} = 0$.
   (e) Otherwise, $\Delta p_{i,t} = 0$.
3. Compute the aggregate price level: $\Delta \ln P_t = \frac{1}{I} \sum_{i=1}^{I} \Delta p_{i,t}$.
4. To compute pass-through coefficients, run the same panel regression as in the empirical analysis on the simulated data.

Supplementary Material

— For Online Publications Only —

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 DQMATCH 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 $c_{j,t}$ 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.

Third, 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.

Fourth, following the recent econometric literature on shift-share designs (Goldsmith-Pinkham, Sorkin, and Swift, 2020; Borusyak, Hull, and Jaravel, 2022), we examine whether the identifying variation in our instrument is driven by the "shares" or the "shifts." We compute the Rotemberg weights associated with our instrument and verify that the largest weights correspond to trade routes where exogeneity of the shipping cost shifts is most plausible (e.g., trans-Pacific routes where congestion was driven by port-level capacity constraints rather than by demand conditions specific to any individual firm or industry).