Delivery Risk, Firm Heterogeneity and Market Power

Manchun Wang*
Boston College

August 2, 2026

Abstract

This paper studies how delivery-time uncertainty affects firms' sourcing networks, market structure, and aggregate prices. Using granular shipment-level data for U.S. importers matched to firm-level financial data, I distinguish realized delivery delays from uncertainty in delivery conditions. I find that firms expand their sourcing networks when delivery uncertainty rises, sourcing from more suppliers, origin countries, and product categories, while realized delays alone do not generate the same response. These sourcing responses are heterogeneous across firms: larger firms adjust more strongly and experience more favorable market-share and markup responses when uncertainty increases.

To quantify the aggregate implications, I develop a multi-sector equilibrium model in which heterogeneous firms choose supplier networks under delivery risk and compete with oligopolistic pricing. Firms differ in sourcing capacity, so delivery uncertainty generates asymmetric cost responses across firms. This heterogeneity reallocates market share toward firms with stronger sourcing positions and creates a markup amplification channel. Feeding the empirical delivery-uncertainty series into the calibrated model, I find that delivery-time uncertainty accounts for about one-fifth of the peak deviation of U.S. manufacturing PPI from its pre-COVID trend, with markup amplification accounting for roughly one-sixth of the model-implied aggregate price increase. The results show that delivery uncertainty affects aggregate prices not only through expected costs, but also by changing market structure and aggregate markups.

Keywords: supplier sourcing, delivery uncertainty, supply chains, market power

JEL Classification: D22, F14, L13, E31

Acknowledgments. The author has no financial relationships or other potential conflicts of interest to disclose.

*PhD candidate at Boston College. Email: wangdup@bc.edu.

1 Introduction

Over the past two decades, global supply chains have been repeatedly exposed to large disruptions that not only lengthen delivery times but also increase uncertainty in delivery performance. Figure 1 plots the ISM Manufacturing Suppliers Delivery Time Index alongside realized delivery risk, defined as the rolling standard deviation of the index over the past 12 months. The figure highlights pronounced episodes of delivery delays associated with major geopolitical, natural, and economic events, including the Iraq War in 2003, the 2011 earthquake and tsunami in Japan, the U.S.China trade tensions in 2018, and the COVID-19 pandemic in 2020. Importantly, these episodes are characterized not only by elevated delivery times but also by heightened volatility and frequent fluctuations in delivery conditions, as reflected in the close comovement between the level index and the delivery risk measure, indicating increased unpredictability in input arrivals.* For firms that rely on timely intermediate inputs, such variability in delivery performance poses a large source of risk. Even when average delivery times decline, greater uncertainty about whether and when inputs will arrive can disrupt production planning and amplify output volatility. As climate-related shocks, geopolitical conflicts, and trade policy uncertainty become more frequent, delivery risk is likely to persist as a central feature of the production environment.

Such delivery risk is costly. alessandria2023aggregate estimate that delays in importing goods lead to a substantial contraction in economic activity, with production, consumption, and employment falling sharply on impact while prices rise. These aggregate effects suggest that delivery frictions can meaningfully disrupt production and pricing.

However, firms may differ in their ability to mitigate delivery risk. Firms with stronger sourcing capabilities may have access to broader sets of potential suppliers, making them more likely to source from lower-cost or more reliable suppliers and to operate more diversified supplier networks. These firms may therefore be more resilient when delivery conditions deteriorate. Survey evidence is consistent with the view that firms actively adjust sourcing strategies in response to supply-chain risk. †

*The ISM Manufacturing Supplier Deliveries Index is a key macroeconomic indicator that measures the pace of delivery of materials from suppliers to manufacturers within the United States. The index is compiled and reported by the Institute for Supply Management (ISM) on a monthly basis, based on surveys of purchasing and supply executives nationwide. High Supplier Deliveries Index suggests that suppliers are facing delays in delivering materials to manufacturers, potentially signaling strong demand and constrained production levels in the manufacturing sector. Conversely, a low index value indicates quick deliveries, which may reflect weaker demand or efficiency improvements in the supply chain.

†For example, a survey by Foley & Lardner reports that, following COVID-19, many manufacturing executives planned to place less emphasis on sourcing from the lowest-cost suppliers and more emphasis on stability and resilience, with a substantial share planning to adopt multi-sourcing strategies. Available online at `https://www.foley.com/wp-content/uploads/imported/9e39e02aa5fa4110b4354cbe3b13cfee.pdf#page=6.13`.

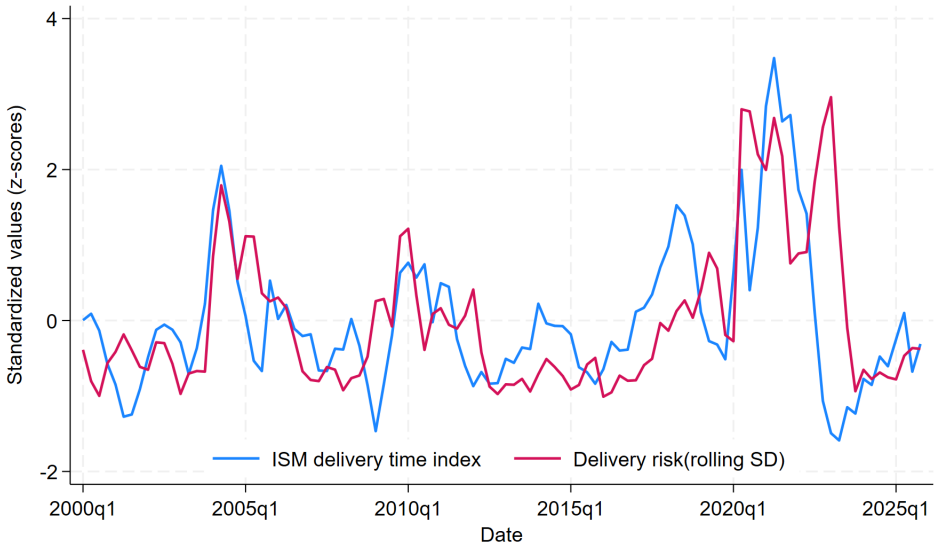


Figure 1: Level and Rolling Standard Deviation of ISM Delivery Time Index

Notes: The figure plots the ISM manufacturing delivery time index and its realized delivery risk, measured as the standard deviation computed over a rolling 12-month window. The series are shown at quarterly frequency from 2000 to 2025, and both the level index and the rolling standard deviation are standardized over the 20002025 sample period. A value of 2 for the delivery risk measure indicates that realized delivery risk is two standard deviations above its historical average.

This heterogeneity in firms' input sourcing positions can have implications beyond the direct increase in production costs. If firms with stronger sourcing positions are better able to absorb delivery disruptions, delivery risk may shift market shares toward these firms. Because firms with larger market shares often charge higher markups under imperfect competition, such reallocation can raise aggregate markups and amplify the effect of delivery risk on aggregate prices.

These observations motivate the question of the paper: how are firms' sourcing networks related to delivery risk, how do these relationships differ across firms, and how does heterogeneity in sourcing positions transmit delivery uncertainty into market structure and aggregate prices? Using granular transaction-level trade data on U.S. importers, I document three sets of empirical facts. First, firms expand their sourcing activity when delivery uncertainty rises: they source from more suppliers, more origin countries, and more product categories, and they increase shipment activity. Second, these responses are heterogeneous across firms. Larger firms are more responsive to delivery uncertainty and expand their sourcing networks more strongly. Third, delivery uncertainty is associated with heterogeneous market outcomes. Larger firms, and firms with broader supplier networks, experience more favorable market-share and markup responses when uncertainty rises. If firms with

larger market shares charge higher markups, this reallocation generates a markup amplification channel that pushes up the aggregate markup and further raises aggregate prices. These findings suggest that delivery risk affects aggregate outcomes not only through a direct cost channel, but also through reallocation across firms.

Motivated by this evidence, I develop a multi-sector equilibrium model in which heterogeneous firms choose supplier networks under delivery risk and compete with oligopolistic pricing. Firms differ in productivity and in their sourcing capacity, so delivery uncertainty affects firms asymmetrically. In the model, higher delivery uncertainty raises expected production costs nonlinearly through the CES production structure, as uncertainty worsens both the expected level and the volatility of material sourcing conditions. Firms with stronger sourcing capacity experience smaller cost increases, gain market share, and charge higher markups. This reallocation pushes up the aggregate markup and amplifies the effect of delivery uncertainty on aggregate prices.

Quantitatively, the model implies that delivery-time uncertainty generated by port congestion can account for a meaningful share of the rise in manufacturing prices during the high-uncertainty period of 2021–2022. Feeding the empirical delivery uncertainty series into the calibrated model, I find that the aggregate price index rises by about 3.2 percent at the peak relative to the pre-COVID uncertainty state. explaining about one-fifth of the peak increase in manufacturing producer prices, while abstracting from aggregate demand, inventory dynamics, energy prices, and other supply-side bottlenecks. Decomposing the model-implied price response shows that higher expected production costs are the primary driver of the aggregate price increase. Because delivery uncertainty raises costs asymmetrically across firms, it also reallocates market shares toward firms with stronger sourcing positions. This reallocation creates a markup amplification channel: market shares shift toward firms with higher markups, raising the aggregate markup and magnifying the effect of delivery uncertainty on aggregate prices. Quantitatively, this markup amplification accounts for roughly one-sixth of the total price increase in the model.

This paper contributes to the growing literature on delivery risk and sourcing networks by linking firm-level sourcing heterogeneity to market-structure and aggregate-price outcomes. Existing work has emphasized how supply-chain frictions affect production costs and quantities. I show that delivery uncertainty can also change the distribution of market power across firms. By combining granular shipment data with a quantitative sourcing model featuring oligopolistic pricing, the paper shows that input-sourcing heterogeneity is an important channel through which delivery risk transmits to market structure. Firms with stronger sourcing positions are better able to absorb delivery uncertainty, so higher uncertainty reallocates market share toward these firms and raises the aggregate markup. This generates a

markup amplification channel beyond the direct expected-cost effect. The quantitative results also show that the effects of delivery uncertainty are nonlinear: when uncertainty is high, expected costs, markups, and aggregate prices increase more rapidly, and the markup amplification becomes larger.

2 Contribution to Literature

This paper contributes to three strands of literature. First, it contributes to the growing literature on supply-chain disruptions and delivery risk. A large body of work studies how disruptions propagate through production networks and affect aggregate outcomes. For example, Carvalho et al. 2021 study the firm-level and aggregate effects of the Thoku earthquake through input-output linkages, while Barrot and Sauvagnat 2016 show that natural disasters affecting upstream suppliers can have persistent effects on downstream firms. More closely related, Alessandria et al. 2023 quantify the aggregate effects of supply-chain delays and depleted inventories during the post-COVID period. This paper builds on these works by distinguishing between realized delivery delays and delivery-time uncertainty. Empirically, I find that firms do not systematically expand their sourcing networks in response to realized delays alone. Instead, sourcing expansion is strongly associated with delivery-time uncertainty: firms source from more suppliers, countries, and product categories when uncertainty rises, and these responses differ systematically across firms. I then focus on how delivery uncertainty affects aggregate prices.

Second, the paper contributes to the literature on endogenous sourcing and input network formation. Models of endogenous supplier choice and production networks show that firms' input choices shape productivity, exposure to shocks, and aggregate outcomes. This includes work on imported input choice and endogenous supplier networks such as Blaum, Lelarge, and Peters 2018, and Dhyne et al. 2023, as well as work on network formation under uncertainty such as Kopytov et al. 2024. This paper incorporates firm heterogeneity in sourcing capacity. Firms with stronger sourcing capacity have access to broader potential supplier pools, which affects both the quality of suppliers they can access and their ability to adjust sourcing when delivery risk increases. This heterogeneity generates asymmetric cost responses across firms.

Third, the paper incorporates market power and output-market reallocation into a supplier-sourcing framework under delivery uncertainty. Firms compete with Atkeson–Burstein pricing, so changes in relative costs affect market shares, residual demand elasticities, and markups. Because delivery uncertainty raises expected costs asymmetrically across firms, input-sourcing heterogeneity is transmitted into output-market structure. Firms with stronger

sourcing positions experience smaller cost increases, gain market share, and charge higher markups. This reallocation raises the aggregate markup and creates an amplification effect: delivery uncertainty affects aggregate prices not only through the direct expected-cost channel, but also through changes in market structure and market power.

The remainder of the paper is organized as follows. Section 2 describes the data and constructs the delivery-risk measures. Section 3 presents empirical evidence on firms' sourcing responses to delivery uncertainty and documents heterogeneity in these responses across firms, as well as the associated market-share and markup responses. Section 4 develops a multi-sector equilibrium model with endogenous supplier sourcing and oligopolistic pricing. Section 5 presents the calibration, baseline quantitative results, and counterfactual exercises on the aggregate price effects of delivery uncertainty.

3 Empirical Analysis

In this section, we describe the data used in the analysis and present empirical evidence on how firms adjust their supplier sourcing in response to delivery risk and the resulting effects on firm outcomes.

First, we document that higher delivery risk induces U.S. importers to diversify their supplier base. Second, beyond pure diversification, we find clear evidence of risk-shifting behavior: U.S. importers reallocate sourcing away from countries with higher delivery risk toward countries with lower delivery risk. Third, we show that delivery-risk-induced diversification reduces firm-level volatility but increases production costs and compresses profitability, a fundamental trade-off between efficiency and stability.

3.1 Data and Variable Construction

3.1.1 Data overview

Panjiva The primary dataset used in this paper is S&P Panjiva, which is a comprehensive bill of lading (BoL) database encompassing more than one billion shipment-level records of cross-border trade transactions. The data includes both import and export transactions for around 20 countries worldwide.[‡] This dataset has become increasingly popular in recent trade and supply-chain research due to its detailed coverage of firm-to-firm international transactions.

A bill of lading is a legal document that serves as evidence that a shipment has been transported from its origin to its final destination. While Panjiva provides shipment-level

[‡]

data for approximately 20 countries, this paper focuses on U.S. imports. For the United States, S&P Global Panjiva obtains bill-of-lading information from U.S. Customs and Border Protection (CBP) through the Freedom of Information Act (FOIA) of 1966. The U.S. BoL data include nearly 200 million records and span the period since 2007.

The bill-of-lading data describe the contractual relationship between the shipper and the consignee. Firms are required to complete a variety of fields in each bill of lading, including the names and addresses of the shipper (exporter) and consignee (importer), descriptions of the goods, vessel name, transport company name, ports of lading (loading) and unlading (unloading), as well as shipment weight and container information. Additionally, the dataset includes supplementary variables imputed by Panjiva itself, such as shipment volume measured in twenty-foot equivalent units (TEUs), Harmonized System (HS) product codes at the six-digit level for each shipment, and - in a limited subset of cases - an estimated transaction value in U.S. dollars. Crucially, the data identify both parties involved in each transaction, namely the consignee (U.S. importer) and the shipper (foreign seller), each of which is assigned a unique company identifier. Panjiva further provides a cross-reference file that allows trade transactions for some shippers and consignees to be linked to their corresponding firms in other datasets, such as S&P Capital IQ, Compustat, and Worldscope. In this paper, we merge the import transaction data with U.S. Compustat to obtain firm-level financial and operational information for importing firms. Table 1 reports the variable names and descriptions for selected key variables in the Panjiva BoL data, including both raw and imputed variables.

Table 1: US import data description for selected variables

Variable	**Description**
panjivarecordid	Unique shipment ID
arrivaldate	Day of arrival
vessel	Name of the vessel that transported the goods
portoflading	Port of lading
portofunlading	Port of unlading
weightkg	Shipment weight in kilograms
volumeteu	Volume of shipment in TEUs
ValueofgoodsUSD	Value of goods in USD
conpanjivaid	Consignee ID
shppanjivaid	Shipper ID
hscode	6-digit HS code
companyid	Capital IQ company ID
shpmtorigin	Location from which shipment left for the US

Congestion Time Index In addition to shipment-level trade data from S&P Panjiva, we use a monthly port-level congestion time index constructed by Bai et al. 2024 to measure delivery delays and supply-chain disruptions. The congestion time index is built using satellite-based data from the Automatic Identification System (AIS) - a vessel tracking system mandated by the International Maritime Organization (IMO) - which records the position, speed, and heading of container ships. This enables measurement of congestion at major ports around the world. The congestion time index measures the average number of hours that a container ship waits in a ports anchorage area before docking at a berth to load or unload cargo. Port congestion time can arise from upstream and downstream frictions along the supply chain, including reduced port-handling efficiency due to health and safety protocols, delays in inland transportation that prevent timely container pickup, and containers remaining unopened at inland facilities because of labor shortages. These frictions propagate through the logistics network and accumulate at ports, ultimately manifesting as prolonged vessel waiting times.

Port congestion waiting time serves as a natural and direct measure of delivery delays for seaborne trade. First, congestion mechanically increases transit time and elevates uncertainty regarding the arrival of imported inputs. Second, containerized maritime shipping plays a central role in international trade - more than half of global trade flows are conducted via maritime shipment - and container ports serve as critical hubs in tightly scheduled logistics networks. When congestion emerges at these hubs, even modest delays can propagate through shipping routes and generate substantial disruptions to downstream production and inventory management. This measure is particularly well-suited to our empirical setting, as U.S. import transactions observed in the Panjiva data are predominantly transported by container ships and therefore necessarily transit through maritime ports.

A key advantage of port congestion time relative to commonly used survey-based delivery measuressuch as the Purchasing Managers' Index (PMI) or Institute for Supply Management (ISM) supplier delivery indicesis that it is relatively more exogenous. Container shipping operates on largely fixed routes under binding service agreements that are predetermined and costly to adjust in the short run. Consequently, port congestion is only weakly influenced by the strategic decisions of individual shippers or by reallocations of capacity across routes in response to demand fluctuations. ibid. argues that, due to this institutional structure, port congestion times are relatively exogenous to demand-side forces that might otherwise confound measures of supply chain disruptions or reflect managers subjective assessments of supply chain conditions.

Finally, port congestion increased sharply and remained elevated during the COVID-19 pandemic, closely aligning with widely documented episodes of global supply-chain stress. This sharp and persistent rise in waiting times provides additional validation that the conges-

tion time index captures economically meaningful delivery delays associated with supply-chain disruptions. The index covers 47 major ports across 24 countries over the period from 2016 to 2025, with these 24 countries accounting for approximately 75 percent of trade transactions and shipments recorded in the Panjiva database.

We use the congestion time index as the primary input to construct our time-varying, firm-level delivery risk measures, which we describe in detail below.

3.1.2 Sample Construction

Our sample construction begins with the complete set of shipments received by U.S. consignees. Initially, we exclude observations lacking panjiva firm identifiers, particularly those with missing *conpanjivaid*. Subsequently, we employ a cross-reference file to connect each *conpanjivaid* with its corresponding *companyid*, which serves as the firm identifier in S&P Capital IQ. Not every consignee can be successfully linked, as numerous firms are privately held and too small for Capital IQ coverage due to scarce publicly available information, even though they participate in international trade. We therefore exclude observations with missing *companyid* from our sample. In our analysis, we work at the level of the ultimate parent US company so we aggregate the data at the ultimate parent company level. We accomplish this by using the S&P Global cross-reference file to associate each companyid with its ultimate parent identifier (*ultimateparentcompanyid*). We remove observations lacking ultimate parent identifiers, though this criterion eliminates only a minimal number of shipments.

Finally, our analysis focuses on the subset of consignees that can be linked to the universe of Compustat firms. We use a crosswalk file obtained from WRDS to merge the Panjiva data with Compustat. This crosswalk provides the start and end dates for each ultimate parent company Capital IQ IDGVKEY linkage, and we merge with the link that is active in a given period. Our sample spans 20162025, corresponding to the availability of the port congestion time index constructed by Bai et al. 2024 and covering both the pre- and post-COVID periods.

Table 2 reports summary statistics for 2019 for two samples: the full sample of U.S. importers with identified ultimate parent Capital IQ firm IDs, and the subsample that can be merged with Compustat, which primarily consists of larger firms. The full sample includes approximately 130,000 ultimate parent firms, while the Compustat-linked sample contains about 7,000 large firms. Across both samples, firm-level import activity is highly right-skewed in terms of the number of shipments, import volume, and the number of shippers. For instance, in the Compustat sample, the average firm conducts roughly 500 import transactions per year, corresponding to about 1,300 containers measured in twenty-foot equivalent units (TEUs). The average firm sources from around 20 distinct shippers annually, whereas

the median firm engages with only four. Firms in the top decile of the distribution import from more than 50 foreign suppliers. The bigger sample exhibits broadly the same patterns as the Compustat public firms sample.

Table 2: Summary Statistics, 2019

Variable	Number of Firms	Mean	Std. Dev.	P10	P50	P90
Full Sample						
Shipments	134,508	56.487	1,470.631	1.000	4.000	78.000
Volume (1000s TEUs)	134,508	140.972	3,121.959	0.100	4.620	147.520
Number of Shippers	134,508	5.233	16.256	1.000	2.000	11.000
Value (million USD)	134,508	10.937	252.812	0.005	0.205	7.246
Compustat						
Shipments	7,014	487.665	6,417.297	1.000	18.000	669.000
Volume (1000s TEUs)	7,014	1,278.534	13,731.060	0.370	26.230	1,523.290
Number of Shippers	7,014	21.771	69.384	1.000	4.000	52.000
Value (million USD)	7,014	116.975	976.069	0.018	1.489	115.504

3.1.3 Construction delivery risk measurement

The congestion time index constructed by Bai et al. 2024 provides a direct measure of delivery delays arising from port congestion and thus captures the first moment of shipping conditions. The central objective of this section, however, is to quantify delivery risk in addition to realized delivery delay, at the firm level. To this end, we construct a shiftshare measure of delivery risk that captures firms exposure to fluctuations in port congestion time. Specifically, we measure firms exposure to shipping time risk using the volatility of the port congestion time index faced by U.S. importers. This firm-level delivery risk measure allows us to examine whether and how firms adopt precautionary sourcing strategies in response to heightened delivery risk. We begin by collapsing the monthly port-level congestion measures to a monthly country-level panel, aggregating across major ports within each country. This aggregation reflects the assumption that congestion at a country's major container ports is representative of overall port congestion conditions faced by shipments entering or leaving that country.

We then propose a three-step approach to construct the delivery risk.

The first step is to isolate unexpected fluctuations in port congestion by residualizing the raw congestion indices. Specifically, we estimate regressions of congestion index on country fixed effects, month-of-year fixed effects, and a linear time trend. Specifically,

$$\text{Congestion}_{ct} = \alpha_c + \lambda_{m(t)} + \beta\, trend + \varepsilon_{ct}, \tag{1}$$

where α_c denotes country fixed effects, $\lambda_{m(t)}$ are month-of-year fixed effects, and *trend* is a linear time trend. Country fixed effects α_c absorb time-invariant differences in port infrastructure, geography, distance from US and baseline congestion levels across countries[§]; month-of-year fixed effects $\lambda_{m(t)}$ capture systematic seasonal patterns in shipping activity; and the time trend removes common global trends in congestion over the sample period. The residual from this regression therefore represents deviations in port congestion that are unexpected given a countrys typical congestion level, seasonal shipping patterns, and long-run global trends. Economically, the residualized congestion index captures high-frequency, within-country innovations to port congestion that are plausibly driven by unanticipated disruptions rather than predictable seasonality or persistent structural differences across ports. This residualization relies on the assumption that the included fixed effects and trend adequately capture the predictable component of congestion, so that the remaining variation reflects unexpected delivery frictions. By focusing on this residual variation, we abstract from slow-moving or systematic factors and instead isolate the component of port congestion most relevant for measuring time-varying delivery risk faced by importing firms.

In the second step, we quantify delivery risk by measuring the volatility of the residualized port congestion index. For each country and month, we compute rolling standard deviations of the residualized congestion measures over six-month and twelve-month windows. These rolling volatilities capture time-varying uncertainty in delivery conditions, abstracting from predictable components of congestion and focusing instead on fluctuations that are unexpected and persistent in maritime logistics.

In the third step, we aggregate country-level congestion volatility to the firm level using firms import sourcing patterns. Specifically, for each firm f and period t, we construct a weighted average of congestion volatility across source countries, where the weights reflect firm f's import shares from each country. In the baseline specification, we use fixed import shares computed prior to 2016, before the start of the congestion index. For each firm, these shares are calculated based on containerized import volumes during the period 2013 - 2015 and are held constant over time. ie.

$$\widehat{\text{DeliveryRisk}}_{f,t}^{(h)} = \sum_c w_{f,c}^{\text{pre}} \sigma_{c,t}^{(h)}, \tag{2}$$

where the weights $w_{f,c}^{\text{pre}}$ are firm f's pre-2016 import shares from source country c, and the country-level standard deviations $\sigma_{c,t}^{(h)}$ capture time-varying volatility in port congestion.

[§]For example, they remove persistent differences across countries in average port congestion, such as structurally longer waiting times in one country relative to another

Weighting country-level congestion volatility by these pre-period shares yields a firm-level delivery risk measure that takes the form of a shiftshare variable: firms that source more heavily from countries experiencing greater volatility in port congestion are more exposed to delivery risk, while firms with more diversified or less exposed sourcing patterns face lower risk. By construction, this approach ensures that firms exposure weights are predetermined with respect to subsequent congestion shocks, isolating variation driven by changes in delivery risk rather than endogenous adjustments in sourcing. Finally, we aggregate the resulting monthly firm-level delivery risk measures to the quarterly frequency by taking within-quarter averages.

3.2 Supplier Sourcing Adjustment Upon Delivery risk

Armed with firm-level measures of delivery risk that capture both the level and volatility of port congestion waiting time, this section documents reduced-form relationships between delivery risk and U.S. importers supplier-sourcing adjustments. Specifically, we estimate the following specification:

$$y_{i,t} = \sum_{\ell=0}^{3} \beta_\ell \, DeliveryRisk_{i,t-\ell} + \mathbf{X}'_{i,t}\gamma + \alpha_i + \delta_{s(i),t} + \varepsilon_{i,t+h} \tag{3}$$

where $y_{i,t+h}$ denotes firm i's supplier-sourcing outcomes at horizon $t+h$, including extensive-margin sourcing adjustment: the number of suppliers, origin countries, HS2 product varieties as well as import activity margins: the number of shipments, and TEU.TEU refers to twenty-foot equivalent units, a standard measure of containerized shipping volume. $DeliveryRisk_{i,t-\ell}$ is the firm-level delivery risk measure, constructed from the level and volatility of port congestion time. We include the contemporaneous value and up to three lags of delivery risk ($\ell = 3$) to allow for dynamic adjustment in firms sourcing decisions. Both the dependent variables and delivery-risk measures are transformed using the inverse hyperbolic sine transformation, so the coefficients can be interpreted approximately as elasticities. $\mathbf{X}_{i,t}$ is a vector of firm-level controls, including firm sales, log total assets, market-to-book ratio, return on assets, and long-term book leverage. α_i denotes firm fixed effects, $\delta_{s(i),t}$ denotes NAICS 4-digit industry-by-time fixed effects, and $\varepsilon_{i,t+h}$ is an error term. The standard error is clustered at firm and time level.

The variable of interest is firm-level delivery risk. We examine how delivery risk affects firms supplier sourcing decisions, capturing both the contemporaneous response and the cumulative effect over time. The shift-share construction mitigates reverse-causality concerns

because firm-level delivery risk is constructed using pre-sample sourcing shares interacted with subsequent country-level port congestion risk. Thus, contemporaneous changes in a firms supply chain network do not mechanically affect its measured exposure. One remaining concern is that this specification may be subject to demand-side endogeneity. We mitigates this concern by including NAICS4-by-quarter fixed effects, which absorb aggregate and industry-specific demand shocks common to firms within the same sector and period. The remaining source of demand endogeneity would therefore have to come from firm-specific demand shocks within an industry-quarter, which we view as a more limited source of bias.

Table 3: Supplier Sourcing Adjustment upon Delivery Risk, Manufacturing

	(1) #Suppliers	(2) HS2 Varieties	(3) #Origin Countries	(4) #Ship	(5) #TEU
Delivery Uncertainty	0.274***	0.208***	0.211***	0.478***	0.580***
	(13.09)	(12.22)	(12.55)	(14.10)	(12.49)
L.Delivery Uncertainty	0.0577***	0.0489***	0.0446**	0.141***	0.159***
	(3.75)	(3.49)	(3.18)	(5.70)	(5.06)
L2.Delivery Uncertainty	0.0441***	0.00594	0.0143	0.0352	0.0558*
	(3.53)	(0.49)	(1.24)	(1.81)	(2.20)
L3.Delivery Uncertainty	0.0662***	0.0355**	0.0426***	0.0966***	0.109***
	(4.62)	(2.60)	(3.69)	(3.95)	(3.33)
Cumulative (L0L3)	0.4423***	0.2984***	0.3128***	0.7513***	0.9029***
SE	0.0376	0.0307	0.0279	0.0646	0.0864
N	21,305	21,305	21,305	21,305	21,305
Adj. Rš	0.843	0.801	0.864	0.851	0.836
Firm FE	yes	yes	yes	yes	yes
NAICS4ŒTime FE	yes	yes	yes	yes	yes

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.010$

Note: This table reports results based on a firm-quarter panel of U.S. manufacturing firms over the 20172025 period. Columns (1)(3) examine the number of suppliers, origin countries, and HS2 product varieties, respectively, while columns (4)(5) examine the number of shipments and TEU volume. Delivery uncertainty is measured as firm-level exposure to the rolling-window volatility of port-congestion waiting times. All regressions control for contemporaneous port congestion waiting time and the one-quarter lags of log sales, log total assets, the book-to-market ratio, return on assets, and long-term book leverage. All specifications include firm fixed effects and four-digit NAICS industry-by-quarter fixed effects. Standard errors are clustered at the firm level.

Table 3 presents reduced-form evidence on how manufacturing importers adjust their sourcing networks in response to second moment delivery risk, measured by the volatility of port congestion waiting time. Our analysis focuses on manufacturing firms, where imported

inputs play a central role in production and supplier-sourcing decisions.

The estimates show that firms exposed to greater delivery uncertainty expand sourcing activity across several margins. In Column (1), higher delivery uncertainty is associated with a significant increase in the number of suppliers, indicating that firms respond to uncertainty by broadening their supplier base. A 10 percent increase in delivery uncertainty is associated with an approximately 2.8 percent increase in the number of suppliers on impact and 4.5% increase over the following four quarters. Columns (2) and (3) show that this diversification is not limited to adding suppliers: firms also expand the range of HS2 product varieties and foreign sourcing countries. These results suggest that importers respond to delivery uncertainty primarily by widening the extensive margin of sourcing. Columns (4) and (5) further show that delivery uncertainty is associated with increases in the number of shipments and total TEU, indicating that firms also adjust shipment frequency and import volume. Overall, the evidence is consistent with a precautionary sourcing mechanism: when exposed to higher delivery uncertainty, firms diversify across suppliers, products, and sourcing countries while also increasing shipment activity to buffer against future disruptions.

Having analyzed firms responses to the volatility of port congestion waiting time, we next turn to its level. We examine whether firms adjust their sourcing behavior in response to realized delivery delays, measured by firm-level exposure to country-level port congestion waiting time. For comparability, we estimate the same specification as in Equation (**??**), replacing the delivery uncertainty measure with a firm level shift-share measure of delivery delays constructed using predetermined pre-sample sourcing weights.

Table 4 examines whether firms adjust their supplier sourcing behavior in response to delivery delays - the first moment of shipping conditions. Across all five sourcing margins - number of suppliers, HS2 varieties, origin countries, shipments, and TEU - we find no evidence that delivery delays induce systematic sourcing adjustments. The contemporaneous coefficients on delivery delay are small, negative, and statistically insignificant, and the same pattern holds for lagged effects up to three periods.

The negative sign of the coefficients is economically intuitive. Realized delivery delays mechanically imply that fewer shipments arrive within a given period, which can temporarily reduce observed shipment counts, supplier relationships, or sourcing variety. In this sense, the first-moment delivery delay captures an ex post realization of disruption rather than an ex ante risk that firms can actively hedge against. This mechanical effect contrasts sharply with the results in Table 3, where firms respond strongly and systematically to delivery uncertainty, measured by the volatility of port congestion time. Together, the results suggest that it is not the level of delivery delays per se that drives sourcing adjustments, but rather the uncertainty surrounding delivery times.

Table 4: Supply Chain Adjustment upon Delivery Time, Manufacturing

	(1)	(2)	(3)	(4)	(5)
	#Suppliers	HS2 Varieties	#Origin Countries	#Ship	#TEU
Delivery Time	-0.0224***	-0.0172**	-0.0152*	-0.0245*	-0.0359**
	(-3.31)	(-2.71)	(-2.58)	(-2.29)	(-2.61)
L.Delivery Time	0.0103	0.00223	0.00698	0.0245**	0.0391***
	(1.80)	(0.42)	(1.43)	(2.89)	(3.36)
L2.Delivery Time	0.00151	0.000968	0.00156	0.00996	0.0148
	(0.28)	(0.18)	(0.31)	(1.20)	(1.48)
L3.Delivery Time	0.00121	-0.00956	-0.00208	-0.0122	-0.0172
	(0.20)	(-1.61)	(-0.41)	(-1.06)	(-1.14)
Cumulative (L0L3)	-0.0094	-0.0235**	-0.0087	-0.0023	0.0008
SE	0.0117	0.0105	0.0097	0.0198	0.0268
N	21,305	21,305	21,305	21,305	21,305
Adj. Rš	0.833	0.794	0.858	0.840	0.826
Firm FE	yes	yes	yes	yes	yes
NAICS4ŒTime FE	yes	yes	yes	yes	yes

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.010$

Note: This table reports results based on a firm-quarter panel of U.S. manufacturing firms over the 20172025 period. Columns (1)(3) examine the number of suppliers, origin countries, and HS2 product varieties, respectively, while columns (4)(5) examine the number of shipments and TEU volume. Delivery time is measured as firm-level exposure to the level of port-congestion waiting times. All regressions control for contemporaneous port congestion waiting time and the one-quarter lags of log sales, log total assets, the book-to-market ratio, return on assets, and long-term book leverage. All specifications include firm fixed effects and four-digit NAICS industry-by-quarter fixed effects. Standard errors are clustered at the firm level.

Finally, we examine whether the estimated response to delivery uncertainty is robust to controlling for the level of delivery delays. To address this concern, we augment the baseline specification by including firms exposure to realized delivery delays and its lags. This allows us to assess whether the effect of delivery uncertainty the second moment of congestion conditionsremains after controlling for changes in the first moment. Table 5 reports the results.

The results show that the precautionary sourcing response of U.S. importers to delivery risk remains strong and largely unchanged. Firms respond to heightened uncertainty in delivery conditions by expanding their sourcing margins.

In contrast, the coefficients on delivery delay levels are negative and often statistically significant, suggesting that realized congestion and shipment delays reduce the volume of shipments and the number of suppliers that firms are able to receive ex post. Taken together, these results underscore that it is the uncertainty surrounding delivery conditions, rather than realized delays themselves, that is associated with firms precautionary sourcing responses. Firms expand sourcing activity ex ante in response to higher delivery risk, even though realized delays may subsequently depress observed import activity.

A key concern is whether the observed increase in suppliers reflects risk-management behavior or simply a scale effect. Since delivery uncertainty is also associated with an increase in HS2 product varieties, firms may add suppliers because they expand product scope or operate at a larger scale, including through love-for-variety forces, rather than because they diversify to insure against delivery risk. To distinguish supplier diversification from product-scope expansion, I re-estimate the sourcing-adjustment specification at the firm-HS2-quarter level:

$$y_{i,h,t} = \sum_{\ell=0}^{3} \beta_\ell DeliveryUncertainty_{i,t-\ell} + \sum_{\ell=0}^{3} \gamma_\ell DeliveryTime_{i,t-\ell} + X_{it}'\theta + \alpha_{ih} + \lambda_{h,t} + \varepsilon_{ih,t},$$

where i indexes firms, h indexes HS2 product categories, and t indexes quarters. The dependent variable ($y_{i,h,t}$) is measured within a given firm-HS2-quarter and includes the number of suppliers, the number of foreign sourcing countries, the number of shipments, and total TEU for firm i in HS2 category h during quarter t. The specification includes firm-HS2 fixed effects, (α_{ih}), which compare the same firm-product category over time, and HS2-by-quarter fixed effects, ($\lambda_{h,t}$), which absorb product-category-specific shocks in each quarter. Similar as before, both dependent variables and delivery risk measurements are transformed using the inverse hyperbolic sine transformation.

Table **??** reports the within-HS2 results. The results show that delivery uncertainty continues to increase the number of suppliers within HS2 categories. This finding suggests

that the supply chain diversification response is not driven solely by firms expanding into new product varieties or by a scale effect associated with broader product scope. Instead, it supports the interpretation that firms diversify their supplier base within existing product categories as a form of insurance against delivery risk.

Another potential concern is whether the observed increase in supply chain expansion reflects an ex-post response to realized delivery volatility increase or a precautionary, forward-looking sourcing strategy. Two features of the data support an expectations-based interpretation. First, both the level and the volatility of port congestion are persistent over time. As shown in Table 7, both country-level congestion volatility and congestion time follows a strongly persistent process, with an autoregressive coefficient around 0.65, implying that shocks to delivery risk decay gradually. Consequently, current congestion condition is informative about near-term future delivery conditions faced by importers sourcing from a given country.

Second, this persistence carries through to firms exposure to delivery risk. Because firm-level delivery risk is constructed as a shift-share weighted average of country-level port congestion measures using predetermined import shares, persistent country-level congestion generates persistent firm-level exposure to both the level and volatility of delivery conditions. Table 7 shows that firm-level delivery risk remains significantly persistent after controlling for firm fixed effects time fixed effects or even industry time fixed effect. These results suggest that current delivery-risk exposure is informative about firms expected near-term risk environment. Therefore, the positive relationship between delivery risk and subsequent sourcing diversification is consistent with a precautionary sourcing response, rather than only an ex-post reaction to realized disruptions.

3.2.1 Firm Heterogeneity in Supplier Sourcing

The previous results show that U.S. importers adjust their sourcing activity in response to delivery risk. However, firms may differ substantially in their ability to make these adjustments. Expanding supplier relationships, sourcing from additional countries, or modifying shipment patterns may require managerial capacity, financial resources, bargaining power, and access to alternative suppliers. These adjustment margins are therefore likely to be less costly for larger firms, which are better able to absorb the fixed costs of searching for and maintaining additional sourcing relationships. To examine this heterogeneity, I augment the empirical specification by interacting firm-level delivery risk with an indicator for large firms:

Table 5: Supply Chain Adjustment upon Delivery Risk, Manufacturing

	(1) #Suppliers	(2) HS2 Varieties	(3) #Origin Countries	(4) #Ship	(5) #TEU
Delivery Uncertainty	0.340***	0.262***	0.260***	0.577***	0.696***
	(14.42)	(13.62)	(13.47)	(10.66)	(13.24)
L.Delivery Uncertainty	0.107***	0.0915***	0.0811***	0.220***	0.248***
	(5.84)	(5.54)	(5.05)	(7.05)	(6.41)
L2.Delivery Uncertainty	0.0714***	0.0291*	0.0338**	0.0818**	0.110***
	(5.03)	(2.18)	(2.63)	(3.50)	(3.79)
L3.Delivery Uncertainty	0.0683***	0.0396**	0.0440***	0.108**	0.122***
	(4.56)	(2.82)	(3.75)	(3.41)	(3.53)
Delivery Time	-0.0522***	-0.0412***	-0.0387***	-0.0765**	-0.0985***
	(-7.26)	(-6.19)	(-6.12)	(-3.20)	(-6.80)
L.Delivery Time	-0.0315***	-0.0302***	-0.0243***	-0.0493*	-0.0484***
	(-4.61)	(-4.82)	(-4.16)	(-2.16)	(-3.38)
L2.Delivery Time	-0.0371***	-0.0265***	-0.0256***	-0.0541**	-0.0621***
	(-5.79)	(-4.44)	(-4.42)	(-2.80)	(-5.05)
L3.Delivery Time	-0.0209**	-0.0224***	-0.0155**	-0.0465*	-0.0576***
	(-2.98)	(-3.36)	(-2.68)	(-2.31)	(-3.34)
Cumulative : Delivery Risk	0.5874***	0.4223***	0.4190***	0.9861***	1.1763***
SE	0.0482	0.0392	0.0361	0.0825	0.1091
Cumulative : Delivery Delay	-0.1416***	-0.1203***	-0.1041***	-0.2264***	-0.2666***
SE	0.0169	0.0147	0.0139	0.0334	0.0380
N	21,305	21,305	21,305	21,305	21,305
Adj. Rš	0.846	0.805	0.866	0.855	0.839
Firm FE	yes	yes	yes	yes	yes
NAICS4ŒTime FE	yes	yes	yes	yes	yes

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.010$

Note: This table reports results based on a firm-quarter panel of U.S. manufacturing firms over the 20172025 period. Columns (1)(3) examine the number of suppliers, origin countries, and HS2 product varieties, respectively, while columns (4)(5) examine the number of shipments and TEU volume. Delivery uncertainty is measured as firm-level exposure to the rolling-window volatility of port-congestion waiting times. Delivery time is measured as firm-level exposure to the level of port-congestion waiting times. All regressions control for contemporaneous port congestion waiting time and the one-quarter lags of log sales, log total assets, the book-to-market ratio, return on assets, and long-term book leverage. All specifications include firm fixed effects and four-digit NAICS industry-by-quarter fixed effects. Standard errors are clustered at the firm level.

Table 6: Supply Chain Adjustment upon Delivery Risk, Manufacturing, within HS2

	(1) #Suppliers	(2) #Origin Countries	(3) #Ship	(4) #TEU
Delivery Uncertainty	0.218***	0.212***	0.409***	0.479***
	(11.11)	(11.89)	(12.10)	(11.75)
L.Delivery Uncertainty	0.0731***	0.0657***	0.132***	0.153***
	(5.08)	(4.64)	(4.17)	(3.90)
L2.Delivery Uncertainty	0.0450***	0.0420**	0.0763**	0.0826*
	(3.75)	(3.10)	(3.02)	(2.34)
L3.Delivery Uncertainty	0.0561***	0.0457***	0.0652*	0.0791
	(4.02)	(3.71)	(2.11)	(2.03)
Delivery Time	-0.0158*	-0.0186*	-0.0352*	-0.0343*
	(-2.06)	(-2.72)	(-2.40)	(-2.17)
L.Delivery Time	-0.0223**	-0.0230***	-0.0377*	-0.0346*
	(-2.93)	(-4.10)	(-2.47)	(-2.06)
L2.Delivery Time	-0.0266**	-0.0254**	-0.0392*	-0.0497*
	(-3.38)	(-3.42)	(-2.37)	(-2.64)
L3.Delivery Time	-0.0137	-0.00929	-0.0298	-0.0268
	(-1.90)	(-1.45)	(-1.97)	(-1.54)
Cumulative : Delivery Risk	0.3923***	0.3658***	0.6820***	0.7941***
SE	0.0354	0.0339	0.0644	0.0794
Cumulative : Delivery Delay	-0.0784***	-0.0762***	-0.1418***	-0.1455***
SE	0.0128	0.0125	0.0238	0.0286
N	121,436	121,436	121,436	121,436
Adj. Rš	0.720	0.740	0.733	0.729
FirmŒHS2 FE	yes	yes	yes	yes
HS2ŒTime FE	yes	yes	yes	yes

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7: Persistence of Delivery Risk(Volatility) and Delivery Delays(Level)

	(1) Country Delivery Delay	(2) Country Delivery Risk	(3) Firm Delivery Delay	(4) Firm Delivery Risk	(5) Firm Delivery Delay	(6) Firm Delivery Risk
L.Country Delivery Time	0.638***					
	(13.24)					
L.Country Delivery Risk		0.682***				
		(11.35)				
L.Firm Delivery Delay			0.558***		0.546***	
			(8.73)		(7.74)	
L.Firm Delivery Risk				0.580***		0.556***
				(14.91)		(12.60)
N	864	864	141,712	141,712	51,893	51,893
Adj. Rš	0.417	0.730	0.579	0.685	0.660	0.723
Firm FE	yes	yes	yes	yes	yes	yes
Time FE	yes	yes	yes	yes		
NAICS4ŒTime FE					yes	yes

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports persistence regressions for delivery risk and delivery delays at the country and firm levels. Delivery risk is measured as a firm-level exposure to the volatility of port congestion waiting times. Delivery delays are measured by firms weighted-average exposure to the level of port congestion time. Columns (1) and (2) present country-level persistence estimates for delivery risk and delivery delay, respectively. Columns (3) and (4) present firm-level persistence estimates for shift-share delivery risk exposure and delivery delay exposure. Firm fixed effects and year-quarter fixed effects are controlled. Standard errors are two-way clustered at the firm and year-quarter levels.

$$y_{i,t} = \beta\, DeliveryUnc_{i,t} + \theta\, Big_i \times DeliveryUnc_{i,t} + Big_i + \mathbf{X}'_{i,t}\gamma + \alpha_i + \delta_{s(i),t} + \varepsilon_{i,t+h} \quad (4)$$

The outcome ($y_{i,t}$) captures firms sourcing adjustments as before, including the number of suppliers, origin countries, HS2 product varieties, shipments, and TEU. Big_i is a time-invariant indicator equal to one if firm i's average annual sales over 2014–2016, measured before the sample period, are above the median among firms in the same four-digit NAICS industry. Since (Big_i) is time-invariant, its main effect is absorbed by firm fixed effects. The coefficient (β) measures the response of non-big firms to delivery risk, while (θ) captures the additional response of large firms.

The outcome $y_{i,t}$ captures firms' sourcing adjustments, including the number of suppliers, origin countries, HS2 product varieties the number of shipments, and TEU. Big_i is a time-invariant indicator equal to one if firm i's average annual sales over 2014–2016, measured before the sample period, are above the median among firms in the same four-digit NAICS industry. Since Big_i is time-invariant, its main effect is absorbed by firm fixed effects. As in the previous specification, both the dependent variables and delivery uncertainty are transformed using the inverse hyperbolic sine, so the coefficients can be interpreted approximately as elasticities. The coefficient β measures the elasticity of non-big firms' sourcing adjust-

ment with respect to delivery uncertainty, while θ captures the additional elasticity for large firms. Therefore, the total elasticity for big firms is $\beta + \theta$.

Table 8: Supplier Sourcing Adjustment upon Delivery Delay, Manufacturing

	(1) #Suppliers	(2) #Countries	(3) HS2	(4) #Ship	(5) TEU
Delivery Uncertainty	0.327***	0.231***	0.253***	0.567***	0.661***
	(10.77)	(10.26)	(10.33)	(11.07)	(10.22)
Big=1 × Delivery Uncertainty	0.0783**	0.0745***	0.0583**	0.143***	0.196***
	(2.55)	(3.17)	(2.29)	(2.67)	(2.85)
SE	22,757	22,757	22,757	22,757	22,757
N	0.849	0.867	0.813	0.857	0.840
Adj. Rš	yes	yes	yes	yes	yes
Firm FE	yes	yes	yes	yes	yes

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.010$

Note: This table reports results based on a firm-quarter panel of U.S. manufacturing firms over the 20172025 period. Columns (1)(3) examine the number of suppliers, origin countries, and HS2 product varieties, respectively, while columns (4)(5) examine the number of shipments and TEU volume. Delivery uncertainty is measured as firm-level exposure to the rolling-window volatility of port-congestion waiting times. Big firms are defined as firms with above-median average annual sales within their four-digit NAICS industry over 2014–2016, before the sample period. All regressions control for contemporaneous port congestion waiting time and the one-quarter lags of log sales, log total assets, the book-to-market ratio, return on assets, and long-term book leverage. All specifications include firm fixed effects and four-digit NAICS industry-by-quarter fixed effects. Standard errors are clustered at the firm level.

Table 8 examines whether large firms adjust their sourcing decisions differently in response to delivery uncertainty. Consistent with the previous baseline results, delivery uncertainty is positively and significantly associated with firms' sourcing and import adjustments across all columns.

Moreover, the interaction term between Big_i and delivery uncertainty is positive and statistically significant across all outcomes. Since Big_i is defined using pre-sample sales and firm fixed effects absorb time-invariant differences across firms, the interaction captures whether firms classified as large respond more strongly to within-firm changes in delivery uncertainty over time. The positive coefficient on interaction term indicates that large firms are more responsive to changes in delivery uncertainty than smaller firms. For column (1), the coefficient on interaction term implies that a 1 percent increase in delivery uncertainty

is associated with an additional 0.1 percent increase in the number of suppliers for big firms relative to non-big firms. Similar patterns appear for origin countries, HS2 product varieties, shipments, and TEU.

These findings are consistent with the model mechanism in which firms diversify their supplier base to hedge against delivery risk. The stronger response among large firms suggests that diversification is not frictionless: firms with greater scale are possibly better able to pay the fixed costs of searching for and maintaining additional supplier relationships. Thus, the empirical results support the interpretation that supplier diversification is an active adjustment margin to delivery risk, and that firm size shapes the ability to use this margin.

This heterogeneity in sourcing adjustment may have broader implications for market structure. If large firms are better able to diversify and maintain import activity when delivery uncertainty rises, they may be better positioned to preserve or expand their market shares relative to smaller firms. Changes in market share can, in turn, affect firms' pricing power and markups. Motivated by this mechanism, I next examine whether delivery uncertainty leads to heterogeneous responses in market share and markup outcomes.

3.3 Market Structure and Markup Implications

3.3.1 Markup

Having documented that firms, especially large firms, expand and diversify their sourcing networks in response to delivery uncertainty, I next examine whether this heterogeneous sourcing capacity translates into differences in market outcomes. If large firms are better able to add suppliers, source from more countries, and maintain input availability when delivery uncertainty rises, they may be less adversely affected than smaller firms and may preserve greater pricing power. To provide direct evidence on this channel, I estimate whether markups respond differently to delivery uncertainty for large and small firms. Firm-level markups are constructed following De Loecker's cost-minimization approach, in which markups are measured as the output elasticity of a variable input divided by that input's revenue share[¶].

Table 9 examines how delivery uncertainty affects firm markups and whether firms with greater sourcing capacity are better able to preserve their market power. Across all specifica-

[¶]Markups are constructed using the production-function approach of De Loecker and Eeckhout(2020), which relies on the firm's cost-minimization first-order condition for a variable input. Specifically, markup is measured as $\mu_{iqt} = \theta^v_{jt}/\alpha^v_{iqt}$, where θ^v_{jt} is the output elasticity of the variable input and α^v_{iqt} is the input's share of firm revenue. In the empirical implementation, I use COGS as the variable input, so $\alpha^v_{iqt} = COGS_{iqt}/Sales_{iqt}$. The benchmark quarterly markup is therefore $\mu_{iqt} = \theta^{COGS}_{jt} \times Sales_{iqt}/COGS_{iqt}$, where θ^{COGS}_{jt} is estimated annually at the two-digit NAICS industry-year level based on De Locker's method and then matched to firm-quarter observations.

tions, the coefficient on delivery uncertainty is negative and statistically significant, indicating that heightened supply-chain uncertainty reduces firm markups. This decline may reflect higher production costs, weaker pricing power, or both.

However, firms with greater sourcing capacity experience significantly smaller declines in markups. Column (1) shows that the interaction between delivery uncertainty and the big-firm indicator is positive and statistically significant, implying that larger firms are better able to maintain their markups when delivery uncertainty rises.

The previous section shows that large firms are better able to diversify their sourcing networks in response to delivery uncertainty. To examine whether sourcing adjustments help explain their greater markup resilience, Columns (2) and (3) interact delivery uncertainty with the number of suppliers and the number of sourcing countries, respectively. Both interaction terms are positive and statistically significant, indicating that firms with more diversified sourcing networks experience smaller declines in markups following an increase in delivery uncertainty. Intuitively, supplier diversification may mitigate supply disruptions by improving sourcing flexibility and reducing dependence on any individual supplier or country, thereby helping firms preserve their pricing power.

Finally, Columns (4) and (5) present horse-race regressions that include both the big-firm interaction and the sourcing-capacity interactions. Once differences in supplier diversification are accounted for, the coefficient on the interaction between the big-firm indicator and delivery uncertainty becomes statistically insignificant, while the interactions with the number of suppliers and sourcing countries remain positive and significant. This pattern suggests that the ability to expand and diversify the supplier network is an important channel through which large firms preserve their markups during periods of heightened delivery uncertainty.

Consistent with the extensive-margin results, Table 10 shows that adjustments along the intensive margin also help firms preserve markups. Firms with more shipments or larger import volumes (TEU) are better able to maintain markups when delivery uncertainty rises, indicating that both expanding sourcing relationships and deepening existing ones contribute to mitigating supply disruptions.

Table 9: Markup Response to Delivery Uncertainty, Manufacturing

	(1) Markup	(2) Markup	(3) Markup	(4) Markup	(5) Markup
Delivery Uncertainty	-0.0156**	-0.0228**	-0.0244**	-0.0283**	-0.0304**
	(-2.12)	(-2.01)	(-2.08)	(-2.36)	(-2.49)
Big=1 × Delivery Uncertainty	0.0118*			0.00717	0.00622
	(1.67)			(1.11)	(0.95)
#Supplier × Delivery Uncertainty		0.00727**		0.00737**	
		(2.01)		(2.18)	
#Origins × Delivery Uncertainty			0.00851**		0.00919**
			(2.13)		(2.45)
SE	22,519	23,124	23,124	22,519	22,519
N	0.900	0.893	0.893	0.900	0.900
Adj. Rš	yes	yes	yes	yes	yes
Firm FE	yes	yes	yes	yes	yes

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.010$

Note: This table reports results from a firm-quarter panel of U.S. manufacturing firms over 20172025. The dependent variable in all columns is the firm-level markup. Delivery uncertainty is measured as firm-level exposure to the rolling-window volatility of port-congestion waiting times. Big firms are defined as those with average annual sales above the median within their four-digit NAICS industry during 20142016, prior to the estimation period. All regressions control for contemporaneous port-congestion waiting times and the one-quarter lags of log sales, log total assets, the book-to-market ratio, return on assets, and long-term book leverage. All specifications include firm fixed effects and four-digit NAICS industry-by-quarter fixed effects. Standard errors are clustered at the firm level.

Table 10: Markup Response to Delivery Uncertainty, Manufacturing

	(1) Markup	(2) Markup	(3) Markup	(4) Markup
Delivery Uncertainty	-0.0254**	-0.0207*	-0.0318**	-0.0278**
	(-2.07)	(-1.84)	(-2.45)	(-2.32)
#Shipment × Delivery Uncertainty	0.00477*		0.00519**	
	(1.95)		(2.22)	
TEU × Delivery Uncertainty		0.00324*		0.00377**
		(1.67)		(2.01)
Big=1 × Delivery Uncertainty			0.00673	0.00723
			(1.05)	(1.12)
SE	23,124	23,124	22,519	22,519
N	0.893	0.893	0.900	0.900
Adj. Rš	yes	yes	yes	yes
Firm FE	yes	yes	yes	yes

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.010$

Note: This table reports results from a firm-quarter panel of U.S. manufacturing firms over 2017 - 2025. The dependent variable in all columns is the firm-level markup. Delivery uncertainty is measured as firm-level exposure to the rolling-window volatility of port-congestion waiting times. Big firms are defined as those with average annual sales above the median within their four-digit NAICS industry during 20142016, prior to the estimation period. All regressions control for contemporaneous port-congestion waiting times and the one-quarter lags of log sales, log total assets, the book-to-market ratio, return on assets, and long-term book leverage. All specifications include firm fixed effects and four-digit NAICS industry-by-quarter fixed effects. Standard errors are clustered at the firm level.

3.3.2 Market Share

Having shown that firms with stronger sourcing capacity are better able to preserve markups under delivery uncertainty, I next examine whether this pricing advantage is accompanied by changes in market share. This exercise is motivated by the standard imperfect-competition logic that firms with larger market shares tend to face less elastic residual demand and therefore have greater ability to sustain higher markups. If sourcing flexibility helps firms avoid production disruptions, stockouts, or dependence on constrained suppliers, these firms may

be better able to maintain sales relative to their competitors when delivery uncertainty rises. The market-share regressions therefore provide evidence on whether sourcing capacity helps firms preserve their relative position within the industry, which is one channel through which sourcing flexibility may translate into higher markups.

Table **??** examines whether firms with greater sourcing capacity are better able to preserve market share when delivery uncertainty rises. The coefficient on delivery uncertainty is generally negative, suggesting that supply-chain uncertainty tends to reduce firms market share. However, the interaction terms with the big-firm indicator, the number of suppliers, and the number of sourcing countries are positive. Column (1) shows that the interaction between the big-firm indicator and delivery uncertainty is positive and statistically significant, indicating that large firms experience a more favorable market-share response than smaller firms.

Columns (2) and (3) show that firms with more suppliers or more sourcing countries also experience smaller market-share declines when delivery uncertainty rises. These results are consistent with the idea that diversified sourcing networks help firms maintain sales relative to competitors during supply disruptions. In Columns (4) and (5), the interaction terms remain positive but become statistically insignificant when firm size and sourcing diversification are included jointly, suggesting that these variables capture overlapping dimensions of sourcing capacity. Overall, the market-share results provide suggestive evidence that sourcing flexibility helps firms preserve their relative position in the industry, although the evidence is weaker than for markups.

Taken together, the empirical results show that delivery risk induces systematic supply-chain adjustments among U.S. importers. Firms respond to higher uncertainty by expanding and diversifying their sourcing networks. These adjustments are heterogeneous across firms: larger firms are more responsive to delivery uncertainty and are better able to expand their supplier networks. This difference in sourcing capacity has important implications for market structure. Firms with stronger sourcing networks experience smaller declines in market share and are better able to preserve markups when delivery uncertainty rises. Moreover, the markup advantage of large firms becomes weaker once supplier diversification is accounted for, suggesting that supplier expansion is an important channel through which large firms maintain pricing power during periods of heightened supply-chain risk.

Table 11: Market Share Response toDelivery Uncertainty, Manufacturing

	(1) Share	(2) Share	(3) Share	(4) Share	(5) Share
Delivery Uncertainty	-0.00445	-0.0136	-0.0187*	-0.0128	-0.0180
	(-0.48)	(-1.33)	(-1.71)	(-1.21)	(-1.61)
Big=1 × Delivery Uncertainty	0.0144*			0.00840	0.00624
	(1.69)			(0.92)	(0.68)
#Supplier × Delivery Uncertainty		0.00863**		0.00543	
		(2.42)		(1.37)	
#Origins × Delivery Uncertainty			0.0104***		0.00770*
			(2.59)		(1.73)
SE	22,549	23,365	23,365	22,695	22,695
N	0.990	0.990	0.990	0.990	0.990
Adj. Rš	yes	yes	yes	yes	yes
Firm FE	yes	yes	yes	yes	yes

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.010$

Note: This table reports results from a firm-quarter panel of U.S. manufacturing firms over 2017 - 2025. The dependent variable in all columns is the firm-level markup. Delivery uncertainty is measured as firm-level exposure to the rolling-window volatility of port-congestion waiting times. Big firms are defined as those with average annual sales above the median within their four-digit NAICS industry during 20142016, prior to the estimation period. All regressions control for contemporaneous port-congestion waiting times and the one-quarter lags of log sales, log total assets, the book-to-market ratio, return on assets, and long-term book leverage. All specifications include firm fixed effects and four-digit NAICS industry-by-quarter fixed effects. Standard errors are clustered at the firm level.

Table 12: Market Share Response to Delivery Uncertainty, Manufacturing

	(1) Share	(2) Share	(3) Share	(4) Share
Delivery Uncertainty	-0.0195*	-0.0167*	-0.0168*	-0.0175
	(-1.87)	(-1.75)	(-1.69)	(-1.64)
#Shipment × Delivery Uncertainty	0.00615***			0.00402
	(2.71)			(1.62)
TEU × Delivery Uncertainty		0.00512***	0.00376*	
		(2.85)	(1.86)	
Big=1 × Delivery Uncertainty			0.00671	0.00743
			(0.73)	(0.82)
SE	23,365	23,443	22,695	22,695
N	0.990	0.990	0.990	0.990
Adj. Rš	yes	yes	yes	yes
Firm FE	yes	yes	yes	yes

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.010$

Note: This table reports results from a firm-quarter panel of U.S. manufacturing firms over 2017 - 2025. The dependent variable in all columns is the firm-level markup. Delivery uncertainty is measured as firm-level exposure to the rolling-window volatility of port-congestion waiting times. Big firms are defined as those with average annual sales above the median within their four-digit NAICS industry during 20142016, prior to the estimation period. All regressions control for contemporaneous port-congestion waiting times and the one-quarter lags of log sales, log total assets, the book-to-market ratio, return on assets, and long-term book leverage. All specifications include firm fixed effects and four-digit NAICS industry-by-quarter fixed effects. Standard errors are clustered at the firm level.

4 Model

To rationalize the empirical evidence on shipping-time risk and quantify its aggregate implications, we develop a theoretical framework that builds on standard models of endogenous sourcing with heterogeneous firms, including Halpern et al. (2015), Gopinath and Neiman (2014), Blaum et al. (2018), and related recent work by Blaum, Esposito, and Heise 2026,Castro-Vincenzi et al. 2024 and Dhyne et al. 2023. As in this literature, establishing

a supplier relationship requires the payment of a fixed cost. Firms may therefore expand their supplier networks in response to greater delivery risk when the diversification benefits of additional suppliers outweigh these fixed costs.

Our framework introduces two key departures. First, firms differ not only in productivity but also in their capacity to search for, establish, and manage supplier relationships, generating heterogeneous sourcing responses to delivery uncertainty. Second, we embed sourcing decisions in an oligopolistic product market with AtkesonBurstein pricing, so that changes in expected sourcing costs affect firms prices, market shares, and endogenous markups. Delivery risk therefore influences aggregate outcomes not only through its direct effect on production costs, but also through the reallocation of market shares and changes in market power across firms.

4.1 Preferences and Demand

Households consume final goods from a finite set of sectors indexed by $k = 1, \ldots, K$. In each sector k, there are N_k final-good producers. Each final-good producer $i = 1, \ldots, N_k$ produces a unique differentiated variety. Preferences have a two-tier CES structure. The upper tier aggregates consumption across sectors, while the lower tier aggregates varieties within each sector.

Aggregate consumption is given by

$$Q = \left(\sum_{k=1}^{K} \alpha_k^{1/\eta} Q_k^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \qquad \sum_{k=1}^{K} \alpha_k = 1,$$

where Q_k denotes sector-k consumption, α_k is the sectoral preference weight, and η is the elasticity of substitution across sectors.

Within each sector, sectoral consumption is a CES aggregate over the differentiated varieties produced by final-good firms:

$$Q_k = \left(\sum_{i=1}^{N_k} q_{ki}^{\frac{\sigma_h-1}{\sigma_h}} \right)^{\frac{\sigma_h}{\sigma_h-1}}, \qquad \sigma_h > 1,$$

where q_{ki} is consumption of variety i in sector k, and σ_h is the elasticity of substitution across varieties within a sector.

Let E denote aggregate final expenditure. The price index associated with the upper-tier

aggregator is

$$P = \left(\sum_{k=1}^{K} \alpha_k P_k^{1-\eta} \right)^{\frac{1}{1-\eta}},$$

where the sectoral price index is

$$P_k = \left(\sum_{i=1}^{N_k} p_{ki}^{1-\sigma_h} \right)^{\frac{1}{1-\sigma_h}}.$$

The CES demand system implies that expenditure on sector k is[‖]

$$E_k = \alpha_k \left(\frac{P_k}{P} \right)^{1-\eta} E.$$

Within sector k, the expenditure share of final-good producer i is

$$s_{ki} = \left(\frac{p_{ki}}{P_k} \right)^{1-\sigma_h} = \frac{p_{ki}^{1-\sigma_h}}{\sum_{j=1}^{N_k} p_{kj}^{1-\sigma_h}}.$$

Therefore, expenditure on variety i in sector k is

$$p_{ki} q_{ki} = s_{ki} E_k = s_{ki} \alpha_k \left(\frac{P_k}{P} \right)^{1-\eta} E.$$

Detailed algebra for demand function and price index see appendix **??**.

4.2 Timing of the Model

The model is static, but decisions are made sequentially within a period. In summary, final-good producers first make the long-run decisions that are costly to adjust: they choose their supplier networks and final-good prices before supplier-specific delivery uncertainty is realized. After delivery states are observed, firms make short-run production decisions, including input sourcing and labor demand.

- **Stage 1: Ex-ante supplier network and pricing.** Each final-good producer forms expectations about the delivery risk associated with its potential suppliers. Taking aggregate expenditure and competitors' prices as given, the firm chooses a supplier

[‖]In the quantitative implementation, sectors are equally weighted: $\alpha_k = 1/K$ for all k. In the model, the cross-sector elasticity η is calibrated very close to one, so the upper-tier CES aggregator is close to the Cobb–Douglas case. When $\eta = 1$, sectoral expenditure shares are fixed at α_k, and $E_k = \alpha_k E$.

network and a final-good price to maximize expected profit. Establishing supplier relationships requires the firm to pay relationship-specific fixed costs. The chosen network determines the set of suppliers from which the firm can source inputs later in the period.

- **Stage 2: Delivery realization, input sourcing, and production.** Supplier-specific delivery states are realized. Conditional on these realizations and on the supplier network chosen in stage 1, the firm chooses input quantities and labor. The firm then produces its final good and sells to consumers at the price chosen ex ante. Profits are realized at the end of the period.

4.3 Intermediate good supplier

Intermediate goods suppliers s use only labor as an input and sell their output to final goods producers. Their production technology is linear:

$$q_s = a_s \ell_s,$$

where a_s is a stochastic, idiosyncratic productivity term drawn independently from a log-normal distribution. The marginal cost for supplier s is given by $\frac{w}{a_s}$, where w is the wage rate. The wage is normalized to one in the quantitative model.

Let p_{si} denote the price charged by supplier s to final good producer i. We assume that supplier s sets this price as a constant markup over marginal cost:

$$p_{si}^M = \mu_{si} \cdot \frac{w}{a_s},$$

where μ_{si} is an exogenous, pair-specific markup. In the baseline model, we set $\mu_{si} = 1$, so that supplier prices equal marginal cost and supplier-buyer relationships generate zero profits. This assumption allows the analysis to focus on final-good producers' sourcing and pricing decisions.

4.4 Final-Good Producers and Delivery Risk

We consider an economy populated by a fixed mass of risk-neutral final-good producers. Producer i in industry k supplies a differentiated manufacturing variety that is sold locally. For notational simplicity, I suppress the industry subscript k when no confusion arises.

Each firm is exogenously matched with a finite set of potential suppliers, denoted by $\mathscr{S}_i$. The size of $\mathscr{S}_i$ varies across firms. Larger and more productive firms are matched with a

larger pool of potential suppliers, capturing their greater ability to search for, screen, and manage supplier relationships. This assumption provides a reduced-form source of heterogeneity in firms' sourcing capacity, consistent with the empirical evidence that larger firms operate broader supplier networks and respond more strongly to delivery risk. Before delivery uncertainty is realized, firm i chooses a supplier network $Z_i \subseteq \mathscr{S}_i$. Only suppliers in Z_i can be used for input sourcing later in the period.

Firm i produces its differentiated variety using labor and a composite of intermediate inputs. Production is given by

$$q_i = \phi_i \left[\alpha_L \ell_i^{\frac{\sigma_{LM}-1}{\sigma_{LM}}} + (1-\alpha_L) M_i^{\frac{\sigma_{LM}-1}{\sigma_{LM}}} \right]^{\frac{\sigma_{LM}}{\sigma_{LM}-1}},$$

where ϕ_i denotes firm productivity, ℓ_i is labor input, and M_i is the intermediate-input composite. The parameter α_L governs the importance of labor in production, while σ_{LM} governs the elasticity of substitution between labor and the intermediate-input composite. In the quantitative model, $\sigma_{LM} < 1$, so labor and intermediate inputs are complements.

The intermediate-input composite is a CES aggregate over the suppliers in the firm's chosen network:

$$M_i = \left[\sum_{s \in Z_i} \alpha_{si}(Z_i)^{\frac{1}{\sigma_M}} m_{si}^{\frac{\sigma_M-1}{\sigma_M}} \right]^{\frac{\sigma_M}{\sigma_M-1}}$$

Here m_{si} denotes the effective quantity of input received by firm i from supplier s, $\alpha_{si}(Z_i)$ is the relative weight of supplier s in firm i's intermediate-input bundle, and σ_M governs the elasticity of substitution across suppliers. Since $\sigma_M > 1$, inputs from different suppliers are substitutes within the intermediate-input composite.

In the benchmark specification, the supplier weights are normalized within the chosen supplier network. Let $a_{si} > 0$ denote the raw supplier-specific input weight. Given Z_i, define

$$\alpha_{si}(Z_i) = \frac{a_{si}^{\sigma_M-1}}{\sum_{r \in Z_i} a_{ri}^{\sigma_M-1}}, \qquad s \in Z_i,$$

so that

$$\sum_{s \in Z_i} \alpha_{si}(Z_i) = 1.$$

This normalization abstracts from mechanical love-of-variety effects in intermediate input use. As a result, expanding the supplier network does not directly raise M_i through a scale effect from adding varieties; instead, the extensive margin of sourcing operates through risk diversification and changes in expected sourcing costs**.

**A perfect-substitutes aggregator would also eliminate mechanical love-of-variety effects. However, it

Delivery Risk A central element of the model is that suppliers are subject to supplier-specific delivery risk. When firm i chooses its supplier network and pays the fixed costs required to establish these relationships, it forms expectations about each supplier's disruption probability ρ_s. At this stage, however, the firm does not know which suppliers will experience a disruption.

Let Γ_s denote the realized delivery fraction of supplier s, defined as the fraction of sourced units that becomes usable input for the buyer. With probability ρ_s, supplier s is disrupted, in which case only a fraction $\underline{\Gamma}_s < 1$ of sourced units is effectively delivered. With probability $1-\rho_s$, supplier s operates normally and all sourced units are delivered. Formally,

$$\Gamma_s = \begin{cases} \underline{\Gamma}_s, & \text{with probability } \rho_s, \\ 1, & \text{with probability } 1-\rho_s, \end{cases} \qquad \text{where } \underline{\Gamma}_s \in (0,1).$$

Delivery disruptions therefore take the form of stochastic iceberg trade costs. If firm i sources x_{si} units from supplier s, only the fraction Γ_s is usable in production, so the effective input received is

$$m_{si} = \Gamma_s x_{si}.$$

Thus, obtaining one effective unit of input from supplier s requires sourcing $1/\Gamma_s$ units, so a lower realization of Γ_s raises the effective cost of usable input[††].

We discipline supplier-level disruption probabilities ρ_s using the empirical relationship between delivery-time uncertainty and realized extreme-delay events. Let u_s denote the delivery-time uncertainty associated with supplier s, and define

$$\rho_s \equiv \Pr(D_s = 1 \mid u_s),$$

where $D_s = 1$ indicates that supplier s experiences a disruption. In the data, a disruption is defined as an event in which realized congestion time exceeds a pre-specified upper-tail cutoff of the empirical congestion-time distribution. Thus, ρ_s captures the probability that a

would imply that inputs from different suppliers are one-for-one substitutes and would typically generate corner input choices. The normalized CES specification instead shuts down the direct scale effect from adding suppliers while preserving imperfect substitutability and a smooth cost function. This keeps the firm's cost and profit objects tractable and substantially simplifies the quantitative solution.

[††]This formulation is a reduced-form way to capture the fact that delays and disruptions reduce the effective input services obtained from a supplier. A lower realization of Γ_s may reflect partial delivery, late arrival, quality deterioration, or costly expediting and reordering. Rather than modeling inventories, backlogs, and shipment timing explicitly, the iceberg formulation summarizes these disruptions as an increase in the effective cost of usable input: obtaining one effective unit from supplier s requires ordering $1/\Gamma_s$ units. This keeps the firm's ex-post cost minimization problem tractable while preserving the model's central mechanism: supplier diversification reduces exposure to idiosyncratic delivery risk.

supplier experiences an extreme delivery delay, conditional on its delivery-time uncertainty. The cutoff choice and the parametric mapping from u_s to ρ_s are described in the calibration section.

4.5 Final-Good Producer's problem

The final good producer's problem is solved by backward induction. After supplier-specific delivery states are realized, final-good producers choose labor and intermediate inputs to minimize production costs. Anticipating these ex-post costs, firms choose their supplier networks and final-good prices before delivery uncertainty is realized.

4.5.1 Ex-post Cost Minimization

Before turning to the firm's ex-ante sourcing and pricing decision, I characterize the unit cost that obtains after delivery states are realized. Given a supplier network Z_i and realized delivery states $\Gamma_i = \{\Gamma_s\}_{s\in Z_i}$, let $c_i(Z_i,\Gamma_i)$ denote the unit cost of producing one unit of the final good. At this stage, the set of available suppliers is fixed, and the firm chooses labor and intermediate inputs taking Z_i and Γ_i as given.

The final-good producer combines labor L_i and the material composite M_i through the CES technology

$$q_i = \phi_i \left[\alpha_L L_i^{\frac{\sigma_{LM}-1}{\sigma_{LM}}} + (1-\alpha_L) M_i^{\frac{\sigma_{LM}-1}{\sigma_{LM}}} \right]^{\frac{\sigma_{LM}}{\sigma_{LM}-1}},$$

where ϕ_i is firm productivity, α_L governs the weight on labor, and σ_{LM} is the elasticity of substitution between labor and the material composite.

The material composite is itself a CES aggregate over the suppliers selected in Z_i:

$$M_i = \left[\sum_{s\in Z_i} \alpha_{si}(Z_i)^{1/\sigma_M} m_{si}^{\frac{\sigma_M-1}{\sigma_M}} \right]^{\frac{\sigma_M}{\sigma_M-1}}.$$

Since delivery disruptions act as iceberg costs, one unit sourced from supplier s generates Γ_s units of usable input. Thus, the effective price of usable input from supplier s is p_s/Γ_s. The dual price index for the material composite is therefore

$$P_{M_i} = \left[\sum_{s\in Z_i} \alpha_{si}(Z_i)\, \Gamma_s^{\sigma_M-1}\, p_s^{1-\sigma_M} \right]^{\frac{1}{1-\sigma_M}} = \Theta_{M_i}^{\frac{1}{1-\sigma_M}}, \tag{5}$$

where $p_s = w/a_s$ is the input price charged by supplier s, and σ_M is the elasticity of substitution across material suppliers. The term

$$\Theta_{M_i} \equiv \sum_{s \in Z_i} \alpha_{si}(Z_i)\, \Gamma_s^{\sigma_M - 1}\, p_s^{1-\sigma_M}. \tag{6}$$

summarizes firm i's realized material sourcing capability. Because $\sigma_M > 1$, Θ_{M_i} is increasing in suppliers' delivery fractions and decreasing in suppliers' input prices. Thus, access to lower-cost or more reliable suppliers raises Θ_{M_i}, while a disruption that lowers Γ_s reduces firm i's realized sourcing capability.

Since the production function is CES and constant return to scale, the marginal cost equals unit cost and given by

$$\begin{aligned} c_i(Z_i, \Gamma_i) &= \frac{1}{\phi_i} \left[\alpha_L^{\sigma_{LM}} w^{1-\sigma_{LM}} + (1-\alpha_L)^{\sigma_{LM}} P_{M_i}^{1-\sigma_{LM}} \right]^{\frac{1}{1-\sigma_{LM}}} \\ &= \frac{1}{\phi_i} \left[\alpha_L^{\sigma_{LM}} w^{1-\sigma_{LM}} + (1-\alpha_L)^{\sigma_{LM}} \Theta_{M_i}^{\frac{1-\sigma_{LM}}{1-\sigma_M}} \right]^{\frac{1}{1-\sigma_{LM}}}. \end{aligned}$$

This expression makes explicit the channel through which sourcing conditions affect marginal cost. Conditional on firm productivity and the wage, the supplier network enters unit cost through the realized material sourcing capability Θ_{M_i}. Lower supplier prices and higher realized delivery fractions raise Θ_{M_i}, lower the material price index P_{M_i}, and reduce unit cost. Conversely, a disruption that lowers Γ_s reduces the effective sourcing capability of the network and raises marginal cost.

4.5.2 Ex-ante sourcing and pricing

Before supplier-specific delivery uncertainty is realized, firm i in sector k simultaneously chooses a supplier network $Z_i \subseteq \mathscr{S}_i$ and a final-good price p_{ki}. The network determines which suppliers are available for input sourcing later in the period. Because the firm observes the distribution of delivery outcomes but not their realizations, delivery risk enters the ex-ante optimization through expected production costs. Firm i therefore chooses its supplier network and price to maximize expected profit.

Different from a monopolistic-competition benchmark, firms compete in an oligopolistic price-setting game. Following Atkeson and Burstein 2008, firms are large within their own sectors: when setting prices, each firm internalizes the effect of its own price on the sectoral price index P_k and hence on its residual demand elasticity. At the same time, firms are small relative to the aggregate economy and take economy-wide objects, such as aggregate expenditure E and the wage w, as given.

Taking competitors' prices $\{p_{kj}\}_{j \neq i}$, aggregate expenditure E, and the wage w as given,

firm i solves

$$\max_{Z_i \subseteq \mathscr{S}_i,\, p_{ki}} \mathbb{E}_{\Gamma_i}\left[\left(p_{ki} - c_i(Z_i,\Gamma_i)\right) q_{ki}(\mathbf{p}_k)\right] - wFC_i(Z_i),$$

where $c_i(Z_i,\Gamma_i)$ is the unit cost associated with supplier network Z_i and delivery-state realization Γ_i, and $FC_i(Z_i)$ is the fixed cost of establishing the supplier network.

Because Z_i and p_{ki} are chosen before the realization of Γ_i, demand is nonrandom conditional on the vector of final-good prices. Delivery uncertainty therefore enters the firm's objective only through realized unit cost. The expected-profit problem can be written as

$$\max_{Z_i \subseteq \mathscr{S}_i,\, p_{ki}} \left(p_{ki} - \mathbb{E}_{\Gamma_i}[c_i(Z_i,\Gamma_i)]\right) q_{ki}(\mathbf{p}_k) - wFC_i(Z_i).$$

For notational convenience, define

$$\bar{c}_i(Z_i) \equiv \mathbb{E}_{\Gamma_i}[c_i(Z_i,\Gamma_i)].$$

Using the demand system derived above, firm i's quantity demanded is[‡‡]

$$q_{ki}(\mathbf{p}_k) = \frac{p_{ki}^{-\sigma_h}}{\sum_{j=1}^{N_k} p_{kj}^{1-\sigma_h}} E_k,$$

where

$$E_k = \alpha_k \left(\frac{P_k}{P}\right)^{1-\eta} E.$$

The first-order condition with respect to p_{ki} yields the Atkeson–Burstein pricing rule. In the Bertrand oligopolistic equilibrium, firm i's price is a markup over its expected unit cost:

$$p_{ki} = \mu_{ki}\bar{c}_i(Z_i), \qquad \mu_{ki} = \frac{\varepsilon_{ki}}{\varepsilon_{ki} - 1}.$$

The residual demand elasticity faced by firm i is

$$\varepsilon_{ki} = \sigma_h(1 - s_{ki}) + \eta s_{ki},$$

where s_{ki} is firm i's market share within sector k, σ_h is the elasticity of substitution across varieties within a sector, and η is the elasticity of substitution across sectors. When $\sigma_h > \eta$, the residual demand elasticity is decreasing in the firm's market share s_{ki}. As a result, larger firms face lower residual demand elasticities and charge higher markups.

This pricing structure links sourcing decisions to market power. Firms with broader

[‡‡] In the Cobb–Douglas upper-tier case, or in the quantitative case where η is close to one, this reduces to $E_k \simeq \alpha_k E$.

sourcing capacity have greater scope to form supplier networks that lower expected unit costs and reduce exposure to supplier-specific delivery disruptions. Lower expected costs allow these firms to set more competitive prices, endogenously expand their market shares, and, through the Atkeson–Burstein markup channel, charge higher markups in equilibrium. This mechanism provides a model-based interpretation of the empirical heterogeneity documented before.

4.5.3 Expected Unit Cost under Delivery Risk

The pricing condition shows how expected unit cost enters the firm's price-setting decision. To understand how supplier-network choices affect $\bar{c}_i(Z_i)$, I now characterize expected unit cost under supplier-specific delivery risk.

The expected unit cost $\bar{c}_i(Z_i)$ depends on the distribution of delivery states through the material sourcing capability Θ_{M_i}. Given the realized unit-cost expression derived above, the supplier network affects expected cost through both the mean and the variance of Θ_{M_i}.

The ex-post unit cost derived above depends on the realized delivery states through the material sourcing capability Θ_{M_i}. Since the supplier network is chosen before delivery states are realized, the firm evaluates each network using the distribution of Θ_{M_i}, and hence the distribution of unit cost.

Given

$$\Theta_{M_i} \equiv \sum_{s \in Z_i} \alpha_{si}(Z_i) \Gamma_s^{\sigma_M - 1} p_s^{1-\sigma_M},$$

and the delivery process is

$$\Gamma_s = \begin{cases} \underline{\Gamma}_s, & \text{with probability } \rho_s, \\ 1, & \text{with probability } 1-\rho_s, \end{cases}$$

the mean of the material sourcing capability is

$$E[\Theta_{M_i}] = \sum_{s \in Z_i} \alpha_{si}(Z_i) p_s^{1-\sigma_M} \left[(1-\rho_s) + \rho_s \underline{\Gamma}_s^{\sigma_M - 1} \right].$$

Assuming supplier-specific delivery shocks are independent across suppliers, its variance is

$$\mathrm{Var}(\Theta_{M_i}) = \sum_{s \in Z_i} \left(\alpha_{si}(Z_i) p_s^{1-\sigma_M} \right)^2 \rho_s (1-\rho_s) \left(1 - \underline{\Gamma}_s^{\sigma_M - 1} \right)^2.$$

Because the outer production aggregator is CES, unit cost is a nonlinear function of Θ_{M_i}:

$$c_i(\Theta_{M_i}) = \frac{1}{\phi_i}\left[\alpha_L^{\sigma_{LM}} w^{1-\sigma_{LM}} + (1-\alpha_L)^{\sigma_{LM}} \Theta_{M_i}^{\frac{1-\sigma_{LM}}{1-\sigma_M}}\right]^{\frac{1}{1-\sigma_{LM}}}.$$

I therefore approximate expected unit cost using a second-order Taylor expansion around $E[\Theta_{M_i}]$:

$$\bar{c}_i(Z_i) \equiv E_{\Gamma_i}[c_i(Z_i,\Gamma_i)] \approx c_i(E[\Theta_{M_i}]) + \frac{1}{2} c_i''(E[\Theta_{M_i}])\mathrm{Var}(\Theta_{M_i}).$$

The full expression for $c_i''(\cdot)$ is provided in Appendix B.1. The mean and variance of the material sourcing term, $E[\Theta_{M_i}]$ and $Var(\Theta_{M_i})$, are computed from the linear representation of Θ_{M_i}, as shown in Appendix B.2. This approximation shows that the supplier network affects ex-ante unit cost through both the mean and the variance of material sourcing capability. A broader supplier network can therefore affect expected cost not only by changing average sourcing conditions, but also by changing exposure to idiosyncratic delivery risk.

Based on the equation of expected cost, delivery uncertainty affects sourcing incentives through both the level and the riskiness of the material input term. Holding the supplier network fixed, an increase in uncertainty raises the probability of a supplier disruption. Since a disrupted supplier has $\Gamma_s < 1$, this lowers $E[\Gamma_s^{\sigma_M - 1}]$ and therefore lowers the mean material sourcing capability, $E[\Theta_{M_i}]$. At the same time, delivery uncertainty changes the dispersion of realized sourcing capability by affecting $Var(\Theta_{M_i})$.

The supplier network determines how exposed the firm is to this risk. With independent or imperfectly correlated supplier shocks, adding suppliers can diversify idiosyncratic delivery risk by spreading exposure across more shocks, thereby reducing $Var(\Theta_{M_i})$. Because supplier weights are normalized within the chosen network, adding suppliers does not generate a mechanical love-for-variety increase in $E[\Theta_{M_i}]$. The value of expanding the supplier network comes from changes in the composition of suppliers and from risk diversification.

4.6 Equilibrium

We characterize a conditional equilibrium taking aggregate final-good expenditure E and the wage w as given, together with a realization of firm- and link-level primitives, including firm productivities, potential supplier sets, link-level productivities, sourcing fixed costs, delivery-risk distributions, and preference and technology parameters. Within each sector k, the equilibrium is a fixed point in firms' sourcing choices, prices, market shares, and

When a firm adds suppliers, the mean $E[\Theta_{M_i}]$ changes only if the added suppliers differ in prices, sourcing weights, or disruption probabilities. By contrast, even identical suppliers can reduce expected cost by lowering $Var(\Theta_{M_i})$, provided their delivery shocks are not perfectly correlated.

markups:

$$\{Z_{ki}, p_{ki}, s_{ki}, \mu_{ki}\}_{i=1}^{N_k}.$$

At this fixed point, each firm's supplier network and price maximize expected profits given competitors' prices and the distribution of supplier delivery risk. Market shares are consistent with the demand system, and markups are consistent with the residual demand elasticities implied by those market shares. Given firms' sourcing decisions and final demand, labor and intermediate-input demands are implied by cost minimization.

The economy-wide conditional equilibrium is the collection of sectoral equilibria for which the resulting sectoral price indexes and sectoral expenditures are consistent with the aggregate price index and aggregate expenditure. The equilibrium determines supplier networks, prices, market shares, markups, quantities, input demands, and profits conditional on E and w. Because the model does not impose a household budget condition or labor-market clearing, aggregate expenditure and the wage are not determined within the model.

5 Quantitative Analysis

5.1 Calibration

5.1.1 Delivery Risk Process

A key input to the quantitative model is the relationship between aggregate delivery-time uncertainty and the probability that a supplier experiences a delivery disruption. I discipline this relationship directly using the data, rather than choosing the disruption probability parametrically.

I normalize delivery-time uncertainty by each origin country's average pre-COVID level to make the measure comparable across countries. Raw congestion volatility may differ systematically across origins because of persistent differences in port infrastructure, shipping composition, or baseline congestion conditions. Under this normalization, $\sigma_{ct}^{d} = 1$ corresponds to country c's normal pre-COVID delivery-time uncertainty. Values above one represent elevated uncertainty relative to that country's baseline, and values below one represent lower-than-normal uncertainty.

This normalization makes cross-country uncertainty comparable and supports the model's common-σ_d specification. In the quantitative model, all suppliers face the same delivery-time uncertainty σ_d, and therefore the same disruption probability. This assumption simplifies the model by replacing a high-dimensional vector of origin-specific uncertainty states with a single risk shifter, allowing the analysis to focus on firms' endogenous sourcing and

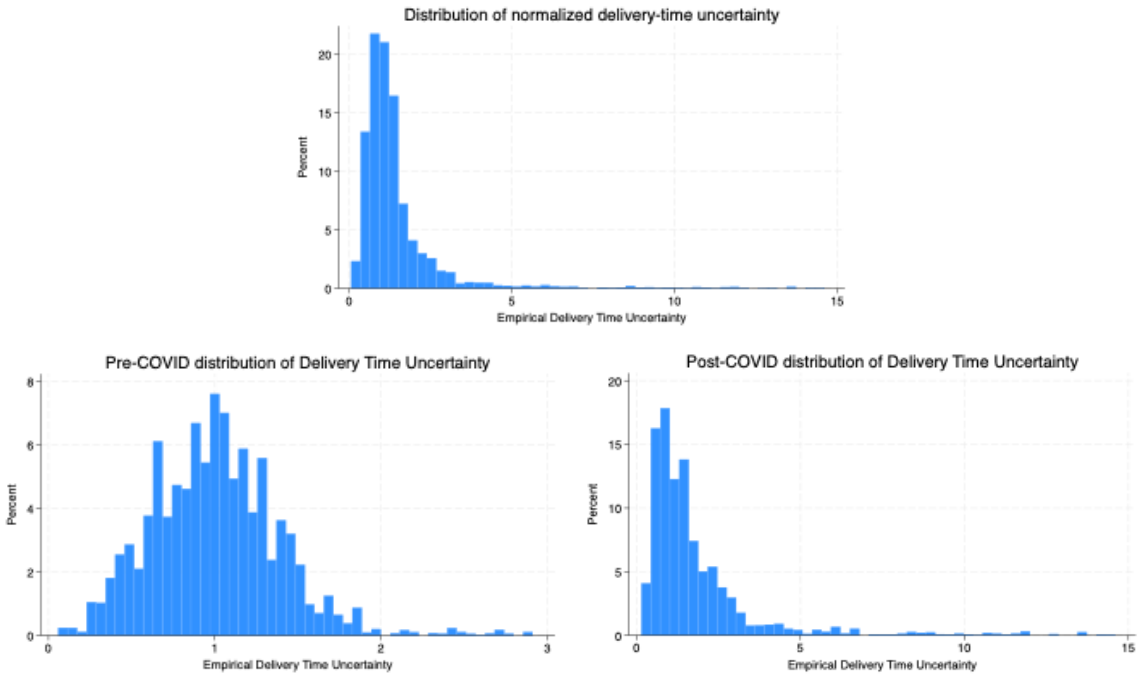


Figure 2: Distribution of Delivery Uncertainty

pricing responses.

Importantly, common ex-ante uncertainty does not imply a common delivery outcome. Conditional on σ_d, supplier-level delivery states remain idiosyncratic and independent across suppliers. Thus, firms face a common ex-ante shipping-risk environment, while their exposure to realized delivery risk depends on the composition and diversification of their supplier networks.

Figure 2 shows the empirical distribution of normalized delivery-time uncertainty. By construction, the pre-COVID distribution(lower-left panel) is centered around one, indicating each country's normal pre-COVID volatility. The post-COVID distribution shifts sharply to the right and becomes substantially more right-skewed, with a long upper tail. This pattern shows that the COVID period was not only associated with a higher average level of delivery-time uncertainty, but also with more frequent extreme uncertainty realizations. The full-sample distribution combines these two regimes and is therefore concentrated near the pre-COVID baseline but displays a long right tail.

We define a disruption as an event in which monthly country-level congestion time exceeds a country-specific pre-COVID percentile cutoff. In the baseline, the cutoff is the 95th percentile of pre-COVID congestion time for each country. We estimate the mapping from normalized delivery-time uncertainty to disruption probability using the shipment-weighted

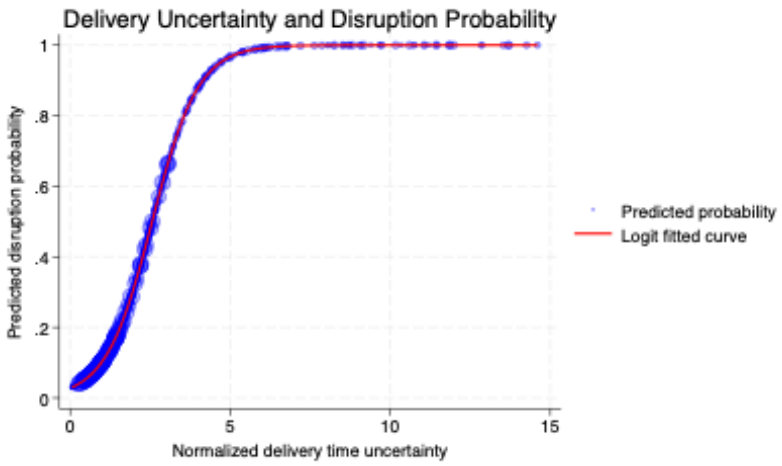


Figure 3: Delivery Uncertainty and Disruption Probability

logit specification

$$\Pr(\text{disruption}_{ct} = 1) = \Lambda\left(a + b\sigma_{d,ct}\right),$$

Here, a_{delay} is the intercept governing the baseline disruption probability, and b_{delay} measures how strongly disruption risk responds to delivery uncertainty. For the baseline 95th-percentile cutoff, the estimates are $a_{\text{delay}} = -3.55$ and $b_{\text{delay}} = 1.39$.

The figure 3 plots the estimated logit mapping from normalized delivery-time uncertainty to disruption probability; the predicted probability remains low near the pre-COVID baseline, $\sigma^d = 1$, but rises rapidly as uncertainty moves above that level.

I construct a shipment-weighted average of the country-level normalized uncertainty measures to summarize the aggregate delivery-risk environment faced by suppliers. I use this aggregate uncertainty series as the empirical input for the delivery-risk states in the quantitative model. Figure 4 plots the average normalized delivery-time uncertainty over time. By construction, the pre-COVID is fluctuated around one due to normalization. The series rises sharply during the COVID period, with the post-COVID average increasing to 1.81, indicating that aggregate delivery-time uncertainty was substantially elevated after 2020.

5.1.2 External Calibration

In the quantitative analysis, we simulate 75 sectors, broadly matching the 77 manufacturing industries represented in the empirical sample. We do not attempt to replicate the full number and size distribution of firms within each empirical industry. Instead, each sector is represented by two heterogeneous final-good producers: one large firm and one small firm. We label the producer with the higher productivity draw as the large firm and the other pro-

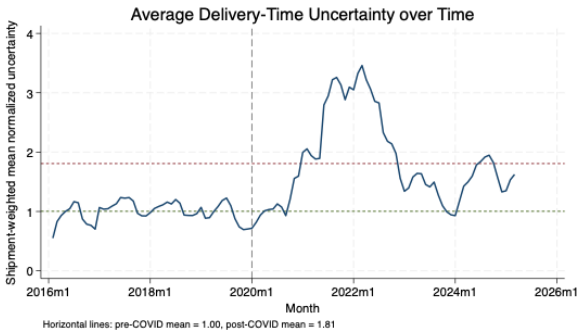


Figure 4: Average Delivery Uncertainty

ducer as the small firm. This abstraction keeps the model computationally tractable while preserving the key source of heterogeneity relevant for our mechanism. Each final good producers should therefore be interpreted as representing one of the two empirical firm-size groups, rather than as an individual firm.

Each final-good producer is randomly matched with a finite set of potential intermediate-input suppliers from which it can source inputs. To capture differences in firms sourcing capacity, large and small firms have access to potential supplier pools of different sizes. We allow a large firm to be randomly matched with up to 100 potential suppliers, compared with 15 for a small firm. These values represent upper bounds on firms sourcing opportunities rather than their realized network sizes. The number of active supplier relationships is determined endogenously by the trade-off between the fixed costs of establishing relationships and the expected cost and diversification benefits of additional suppliers.

This heterogeneity in potential supplier-pool size is motivated by the substantial difference in supplier-network size observed in the manufacturing sample. Based on the empirical moments, small firms source from approximately 6.79 suppliers on average, whereas large firms source from approximately 23.98 suppliers.

The simulated economy contains approximately 3,000 intermediate-good suppliers, implying a supplier-to-final-good-producer ratio of about 20 to 1. Intermediate suppliers may serve multiple buyers, so the large supplier universe provides substantial scope for heterogeneous supplier matches without mechanically constraining firms network choices.

Because the model characterizes an equilibrium conditional on the wage and aggregate expenditure, we normalize both to one in the quantitative analysis: $w = 1$ and $E = 1$.

In the model, final-good producers and intermediate-good suppliers draw their productivities independently from a common log-normal distribution. We normalize the location parameter of log productivity and discipline firm heterogeneity primarily through its dispersion parameter. The location parameter largely determines the units in which productivity

and production costs are measured. Because the wage and aggregate expenditure are normalized and the quantitative analysis focuses on relative outcomes, the absolute level of productivity is not central to the models mechanism. In contrast, productivity dispersion governs the relative productivity gap between firms. It therefore affects which producer is classified as large within each sector, the resulting differences in marginal costs and market shares, and firms' incentives to establish supplier relationships. We calibrate the productivity-dispersion parameter to match the difference in market shares between large and small firms, as described below.

We set the elasticity of substitution across final-good varieties to $\sigma_h = 4.5$. Estimates of this demand elasticity typically range between 3 and 6 de2011product; blaum2018gains, for example, use $\sigma_h = 3.83$. We set the elasticity of substitution across sectors to $\eta = 1.15$, in line with values used in the literature, including 1.19 in grossman2024tariffs, 1.24 in edmond2015competition, and 1.01 in atkeson2008pricing. The within-sector elasticity is higher than the across-sector elasticity, so that firms with larger within-sector market shares face less elastic residual demand and can charge higher markups. The labor share in the production function is set to $\alpha_L = 0.5$. The fixed sourcing cost governs the cost of expanding the supplier network and, in models with a love-of-variety motive, can be disciplined by the average number of suppliers. However, in the model without love for variety, firms therefore have little incentive to multi-source in the no-uncertainty economy, and baseline sourcing is often at a one-supplier corner. As a result, the fixed sourcing cost is weakly identified from average supplier counts. I therefore fix the sourcing-cost scale at a baseline value and treat it as an auxiliary parameter.

Table 13 summarizes those exogenous primitives and externally calibrated parameters.

5.1.3 Internal Calibration

After fixing the externally calibrated parameters, we jointly calibrate four parameters by simulated method of moments: the dispersion of firm productivity Φ^{disp}, the elasticity of substitution across material suppliers σ_M, the elasticity of substitution between labor and materials σ_{LM}, and the severity of the idiosyncratic delivery shock x_{idio}.

Productivity dispersion is primarily disciplined by the difference in market shares between large and small firms, since greater dispersion generates larger productivity, marginal-cost, and market-share differences within sectors. The elasticity σ_M is mainly informed by the difference in the diversification elasticity between large and small firms, as it governs how effectively firms can substitute across suppliers and therefore the value of expanding the supplier network when delivery risk increases.

The parameters $\underline{\Gamma}$ and σ_{LM} are jointly disciplined by the responses of market shares and

Table 13: Externally Calibrated Parameters

Parameter / Object	Notation	Value	Source / Discipline
Number of sectors	K	75	Broadly matches the 77 manufacturing industries in the empirical sample.
Potential suppliers for large firms	$\bar{K}_B$	100	Upper bound on sourcing opportunities for large firms;
Potential suppliers for small firms	$\bar{K}_S$	15	Upper bound on sourcing opportunities for small firms;
Intermediate-good suppliers	M	3,000	
Wage	w	1	Numeraire normalization.
Aggregate expenditure	E	1	Numeraire normalization.
Elasticity across final-good varieties	σ_h	4.5	Within the range of estimates in de2011product; blaum2018gains use $\sigma_h = 3.83$.
Elasticity across sectors	η	1.15	grossman2024tariffs, edmond2015competition, and atkeson2008pricing.
Labor share	α_L	0.5	Baseline production-function parameter.

markups to delivery uncertainty interacted with the number of supplier links. The parameter $\underline{\Gamma}$ governs the severity of the material-input loss conditional on a disruption, while σ_{LM} determines how strongly changes in material sourcing costs pass through to total marginal cost, final-good prices, and market shares. Through the Atkeson–Burstein pricing channel, the resulting changes in market shares then translate into changes in markups.

Table 14 reports the internally calibrated parameters and the model fit. The model closely matches the targeted moments, including the big-small gap in the diversification response to delivery uncertainty and the coefficients on the interaction between uncertainty and supplier links in the markup and market-share regressions. The model also matches the average within-sector market shares of large and small firms.

Beyond the targeted moments, the model performs reasonably well on untargeted validation moments. In particular, it matches the baseline coefficient on delivery uncertainty in the market-share regression with supplier-link interactions, as well as the big-small gap in the market-share response to uncertainty. It also generates roughly half of the observed big-small gap in the markup response to uncertainty, even though this moment is not directly targeted.

In the markup and market-share regressions, we target only the interaction between de-

Table 14: Internally Calibrated Parameters and Moment Fit

Panel A. Internally calibrated parameters		
Parameter	Notation	Value
Productivity dispersion	Φ^{disp}	1.40
Elasticity across material suppliers	σ_M	10.00
Elasticity between labor and materials	σ_{LM}	0.35
Idiosyncratic delivery-loss severity	$\underline{\Gamma}_s$	0.82
Panel B. Targeted moments		
Moment	Model	Data
Average within-sector market share, large firms	87	89.3
Average within-sector market share, small firms	13	10.7
Big-small gap in diversification response to uncertainty	0.077	0.078
Markup coefficient on uncertainty $\times$ supplier links	0.0099	0.0073
Market share coefficient on uncertainty $\times$ supplier links	0.0098	0.0086
Panel C. Untargeted moments		
Moment	Model	Data
Standard deviation of log sales	2.19	2.360
Market-share coefficient on uncertainty in link-interaction regression	-0.0142	-0.0136
Big-small gap in market-share response to uncertainty	0.0132	0.0144
Big-small gap in markup response to uncertainty	0.0050	0.0118

livery uncertainty and the number of supplier links, rather than the baseline coefficient on delivery uncertainty. The interaction term more directly captures the mechanism emphasized by the modelhow sourcing capacity shapes firms responses to delivery risk. By contrast, the baseline coefficient measures the average response to uncertainty and may also reflect demand shocks, aggregate conditions, pass-through, and other adjustment channels that are not explicitly included in the model. We therefore leave the baseline coefficient untargeted to avoid forcing the model to account for mechanisms outside its scope.

We also do not target the difference between large and small firms in the elasticity of markups or market shares with respect to uncertainty. Firm size is associated with many dimensions of heterogeneity beyond sourcing capacity, such as demand, financing conditions, inventory management, and organizational structure. The observed largesmall gap may therefore reflect several mechanisms that are not explicitly modeled. Instead, we target the interaction between delivery uncertainty and the number of supplier links, because this moment more directly isolates how supply-chain capacity shapes the transmission of delivery risk into market shares and markups.

$\sigma_{LM} = 0.35$ means that labor and material are complement instead of substitute. In the model without love for variety, delivery risk affects firms primarily through disruptions to the material composite. The elasticity between labor and materials then becomes much more important, because it determines how much those material disruptions translate into total marginal cost. When σ_{LM} is low, labor and materials are complements, so disruptions to the material composite cannot be easily offset by labor substitution and have large cost effects. When σ_{LM} is high, firms can substitute away from disrupted materials, substantially dampening the effect of delivery risk. Thus, in the model without love for variety, a relatively low σ_{LM} is needed for sourcing risk to matter quantitatively.

5.2 Baseline Simulation Result

Figure 5 shows how the calibrated economy responds to higher delivery-time uncertainty. In the model without love for variety, firms have little incentive to source from multiple suppliers when uncertainty is low, since the cost-minimizing choice is to concentrate purchases with the most attractive supplier. For this reason, the level of active supplier links is close to one in low-uncertainty states and should not be interpreted literally as the empirical number of suppliers. The relevant object is instead the responsiveness of sourcing choices to uncertainty. As delivery uncertainty rises, firms expand their supplier networks, and large firms are more responsive along this margin because they have broader potential sourcing opportunities.

This sourcing-capacity heterogeneity generates heterogeneous firm-level responses and induces reallocation across firms. The within-sector market-share and markup panels in Figure 5 summarize the reallocation channel. As delivery uncertainty rises, large firms gain market share while small firms lose market share. This pattern reflects two effects of having a larger potential supplier pool. First, a larger pool is more likely to include highly productive suppliers, so large firms can access better suppliers than small firms and face lower material-input costs. This selection effect is a first-order sourcing-capacity advantage. Appendix **??** shows analytically that higher supplier productivity will raise the material sourcing term Θ_i, lower the firms material cost share, and therefore reduces the sensitivity of unit cost to delivery uncertainty. Second, a larger pool gives large firms more room to expand their active supplier networks when delivery risk increases. Multi-sourcing reduces the volatility of realized input delivery through diversification, which lowers expected costs by limiting the cost consequences of supplier-level disruptions. As large firms relative costs rise less, market share shifts toward them. Under Atkeson–Burstein pricing, this market-share reallocation reduces large firms residual demand elasticities and raises their markups, while small firms

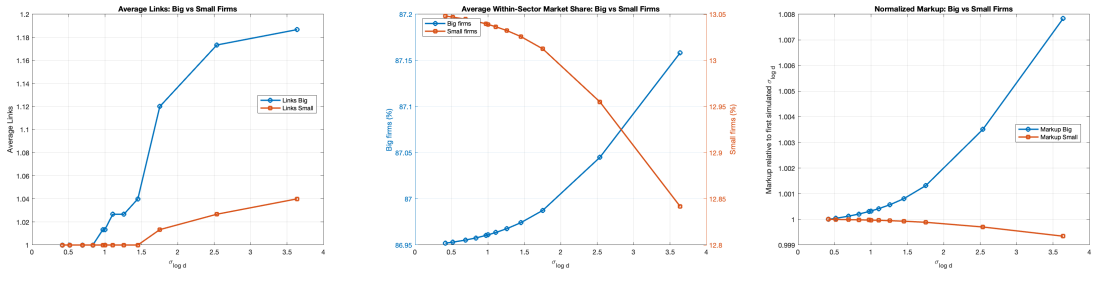


Figure 5: Heterogeneous firm responses to delivery-time uncertainty

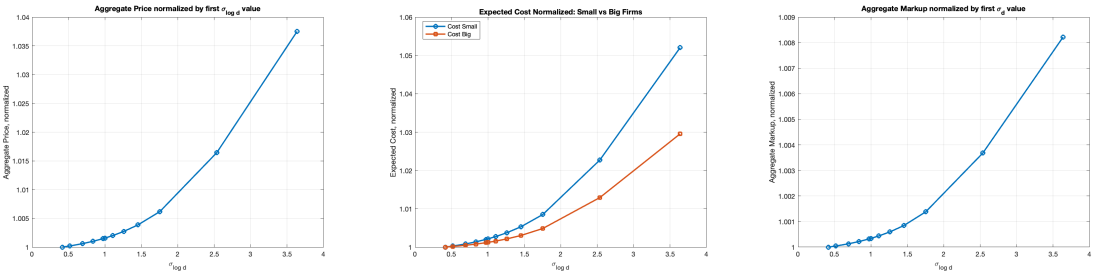


Figure 6: Aggregate price, expected cost, and markup

lose market share, face more elastic residual demand, and slightly lower their markups.

Figure 6 summarizes the aggregate implications of higher delivery-time uncertainty. The aggregate price index rises monotonically with uncertainty, reflecting both higher expected production costs and changes in market power. The simple average expected-cost panel shows that costs increase for both large and small firms, but the increase is larger for small firms. This difference reflects their weaker sourcing position: small firms have fewer potential sourcing options, so disruptions to material inputs translate more strongly into marginal costs. Large firms expected costs also rise, but by less, because their broader supplier access allows them to rely on better suppliers and adjust sourcing more effectively when risk increases. Moreover, the cost response is nonlinear: when delivery uncertainty is already high, a further increase generates a larger rise in expected cost than the same increase would in a low-uncertainty environment.

The sales-weighted aggregate markup also rises with delivery-time uncertainty as market shares shift toward larger firms, which charge higher markups. This reallocation further amplifies the aggregate price response, particularly at higher levels of uncertainty, when the cost gap between large and small firms becomes more pronounced. Overall, the figure shows that delivery uncertainty raises aggregate prices primarily through higher expected production costs, with an additional amplification arising from the reallocation of market shares toward larger, higher-markup firms.

5.3 Counterfactual Analysis

Armed with the calibrated model, we conduct several counterfactual analyses, including a time-series decomposition of the aggregate price response, a simulation that eliminates sourcing-capacity differences between large and small firms, and a simulation in which supplier networks are held fixed to shut down endogenous sourcing adjustment.

5.3.1 Time series Decomposition

To assess the quantitative importance of delivery-time uncertainty over the observed period, I next feed the empirical uncertainty series into the calibrated model and decompose the resulting aggregate price response into cost and markup components.

Figure **??** feeds the empirical To assess the quantitative importance of delivery-time uncertainty over the observed period, I next feed the empirical uncertainty series into the calibrated model and decompose the resulting aggregate price response into cost and markup components.

For each value of delivery-time uncertainty, firm j's price can be written as

$$p_j(\sigma_t) = \mu_j(\sigma_t) c_j(\sigma_t),$$

where $\mu_j(\sigma_t)$ is the firm's markup and $c_j(\sigma_t)$ is its expected marginal cost. Since the aggregate price index is nested CES, its log change can be written as a share-weighted average of firm-level price changes. For two adjacent uncertainty states σ_k and σ_{k+1}, I therefore decompose the aggregate price change as

$$\Delta \log P_k = \underbrace{\sum_j \bar{s}_{jk} \left[\log \mu_j(\sigma_{k+1}) - \log \mu_j(\sigma_k)\right]}_{\text{Markup contribution}} + \underbrace{\sum_j \bar{s}_{jk} \left[\log c_j(\sigma_{k+1}) - \log c_j(\sigma_k)\right]}_{\text{Cost contribution}} + R_k.$$

where

$$\bar{s}_{jk} = \frac{1}{2}\left[s_j(\sigma_k) + s_j(\sigma_{k+1})\right]$$

is the average nested-CES expenditure share across the two equilibria, and R_k is a finite-change approximation residual.

Figure 7 feeds the empirical shipment-weighted average delivery-time uncertainty series, shown in 4, into the calibrated model and computes the corresponding sequence of equilibria. For each uncertainty state, the model solves for firms' sourcing choices, prices, market shares, markups, and expected marginal costs. The figure then decomposes the model-implied change in the aggregate price index from the pre-COVID uncertainty state, $\sigma = 1$,

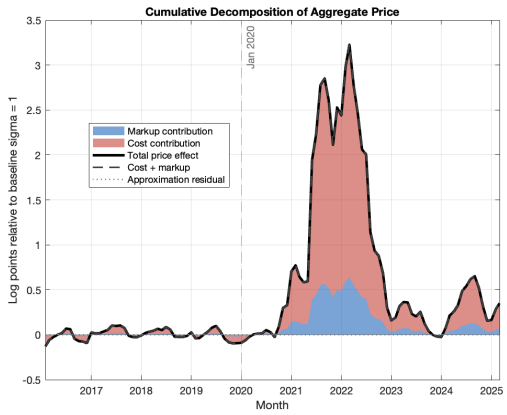


Figure 7: Time-series decomposition of the model-implied aggregate price response



Figure 8: Manufacturing producer prices, BLS

into expected-cost and markup components. The blue area in the figure plots the cumulative markup contribution,

$$\sum_{\ell \le k} \sum_{j} \bar{s}_{j\ell} \left[\log \mu_j(\sigma_{\ell+1}) - \log \mu_j(\sigma_\ell) \right],$$

while the red area plots the cumulative expected-cost contribution,

$$\sum_{\ell \le k} \sum_{j} \bar{s}_{j\ell} \left[\log c_j(\sigma_{\ell+1}) - \log c_j(\sigma_\ell) \right].$$

The decomposition result shows that the aggregate price response is driven primarily by higher expected production costs, while the markup channel provides a meaningful amplification, especially when delivery-time uncertainty is high. During the high-uncertainty episode of 2021–22, the aggregate price level rises by approximately 3.2 percent at its peak, with the markup channel accounting for roughly one-sixth of the total increase.

To gauge the magnitude of the model-implied effect, I compare the simulated price response with the deviation of the BLS manufacturing PPI from its pre-COVID log-linear

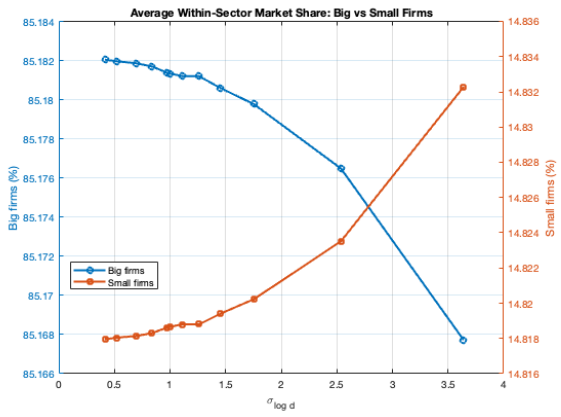


Figure 9: Firm responses without sourcing-capacity heterogeneity

trend, shown in Figure 8. The model implies a peak increase in the aggregate price index of about 3.2 log points, while the detrended manufacturing PPI rises by roughly 15 percent at its peak. Taken literally, delivery-time uncertainty due to port congestion therefore accounts for about one-fifth of the peak deviation of manufacturing producer prices from their pre-COVID trend, even though the model abstracts from aggregate demand, inventory dynamics, supplier capcacity constraint and other supply-side bottlenecks.

5.3.2 The Role of Sourcing Heterogeneity

Next, I examine how the results change when large and small firms have the same sourcing capacity. Specifically, I set the maximum size of the potential supplier pool to 100 for both firm, eliminating large firms advantage in accessing better suppliers and their greater flexibility to diversify across suppliers.. Figure 9 shows that, as delivery uncertainty increases, the within-sector market share of large firms declines slightly, while the market share of small firms increases. Thus, shutting down sourcing-capacity heterogeneity reverses the direction of the market-share reallocation relative to the baseline, although the magnitude of the reallocation is small.

Figure 10 shows that, in this counterfactual, the aggregate price response is almost entirely driven by higher expected costs. Because market share no longer reallocates toward large firms, the markup channel does not provide meaningful amplification.

5.3.3 The Role of Endogenous Sourcing

Figure 11 compares the endogenous-sourcing economy with a counterfactual in which firms' supplier networks are fixed at their no-uncertainty choices. The comparison isolates the incremental role of sourcing adjustment after the initial supplier network has already been

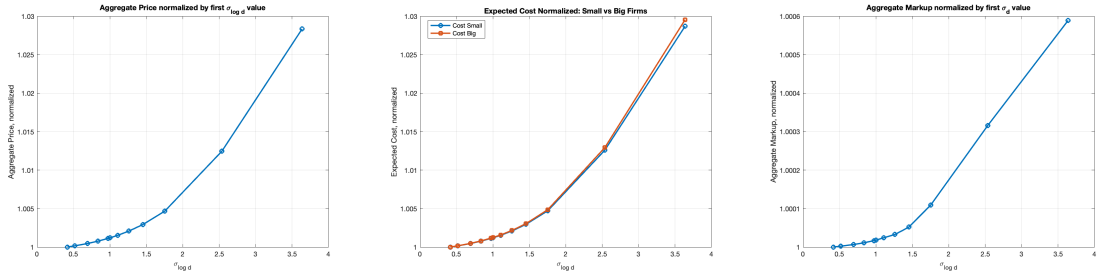


Figure 10: Aggregate responses without sourcing-capacity heterogeneity

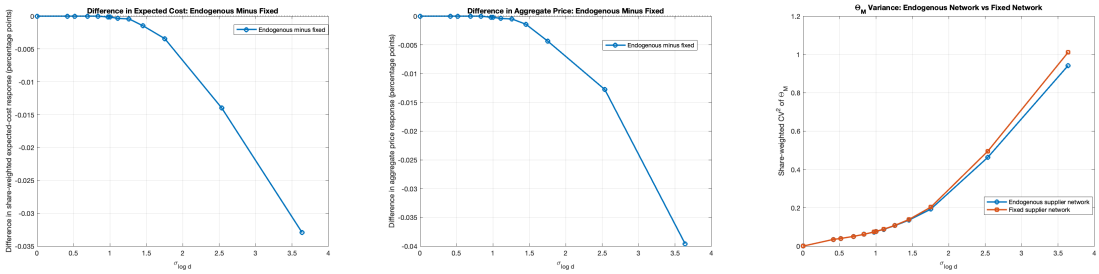


Figure 11: Endogenous Sourcing Adjustment Relative to Fixed Supplier Networks

chosen.

The figure shows that endogenous sourcing reduces the variance of the material sourcing term and slightly lowers expected marginal costs and aggregate prices relative to the fixed-network counterfactual. The effect is economically intuitive but quantitatively small. The reason is that the fixed-network benchmark is not a random or inefficient sourcing allocation: it fixes the network chosen by optimizing firms in the no-uncertainty economy. Firms have already selected suppliers that deliver high values of the material sourcing term Θ_M. When uncertainty rises, further network adjustment can improve Θ_M or reduce risk exposure, but the marginal gain from reoptimization is limited.

6 Conclusion

This paper studies how firms adjust their supply chains in response to delivery risk and how heterogeneity in sourcing networks affects market structure and aggregate prices. Using granular shipment-level data from S&P Panjiva matched to firm-level financial data, I document that U.S. importers expand their sourcing networks when delivery uncertainty rises.

The results imply that the dominant source of reallocation and markup amplification comes from heterogeneity in firms supplier quality. In particular, differences in sourcing capacity allow large firms to access higher-productivity suppliers, generating supplier-quality differences that shape firms cost responses to delivery uncertainty. Diversification instead acts as a mitigating margin: it dampens the cost and price effects of delivery uncertainty, with somewhat larger cost mitigation for large firms, but its quantitative effect is small.

This adjustment is heterogeneous: larger firms respond more strongly by adding suppliers, sourcing from more countries, and expanding import activity. I also show that larger firms experience more favorable market-share and markup responses when delivery uncertainty increases, and that these differences are partly associated with their broader sourcing networks.

To quantify the aggregate implications, I develop a multi-sector equilibrium model in which heterogeneous firms choose supplier networks under delivery risk and compete using Atkeson–Burstein pricing. Firms differ in productivity and in sourcing capacity, where sourcing capacity is modeled as access to a larger potential supplier pool. This heterogeneity generates reallocation when delivery uncertainty rises. Firms with broader sourcing opportunities are more likely to access high-productivity suppliers and have greater flexibility to diversify across suppliers. As a result, their expected costs rise less under delivery uncertainty, allowing them to gain market share. Through the Atkeson–Burstein pricing channel, this reallocation toward larger firms raises aggregate markups and amplifies the aggregate price response beyond the direct expected-cost channel.

The quantitative results imply that delivery-time uncertainty raises aggregate manufacturing prices substantially during the high-uncertainty episode around 2021–2022. The response is nonlinear: both expected costs and markup amplification become stronger when delivery uncertainty is high. At the peak, the model-implied aggregate price increase is about 3.2 % relative to the pre-COVID uncertainty state, with the markup channel accounting for roughly one-sixth of the total response. Compared with the deviation of U.S. manufacturing PPI from its pre-COVID trend, the model suggests that delivery-time uncertainty can account for about one-fifth of the peak increase in manufacturing producer prices, while abstracting from aggregate demand, inventory dynamics, energy prices, and other supply-side bottlenecks.

Overall, the paper highlights two mechanisms through which delivery uncertainty matters. First, uncertainty raises expected production costs nonlinearly. Second, heterogeneity in input sourcing positions transmits delivery risk into output-market reallocation and markup amplification. Delivery risk therefore affects aggregate prices not only by raising firms' marginal costs, but also by changing market structure and the distribution of market power across firms.

References

Acemoglu, Daron and Pablo D Azar (2020). "Endogenous production networks". In: *Econometrica* 88.1, pp. 33–82.

Ahir, Hites, Nicholas Bloom, and Davide Furceri (2022). *The world uncertainty index*. Tech. rep. National bureau of economic research.

Alessandria, George et al. (2023). "The aggregate effects of global and local supply chain disruptions: 2020–2022". In: *Journal of international Economics* 146, p. 103788.

Antras, Pol, Teresa C Fort, and Felix Tintelnot (2017). "The margins of global sourcing: Theory and evidence from US firms". In: *American Economic Review* 107.9, pp. 2514–2564.

Atkeson, Andrew and Ariel Burstein (2008). "Pricing-to-market, trade costs, and international relative prices". In: *American Economic Review* 98.5, pp. 1998–2031.

Bachmann, Rüdiger and Christian Bayer (2013). "Wait-and-Seebusiness cycles?" In: *Journal of Monetary Economics* 60.6, pp. 704–719.

Bai, Xiwen et al. (2024). "The causal effects of global supply chain disruptions on macroeconomic outcomes: evidence and theory". In.

Barrot, Jean-Noël and Julien Sauvagnat (2016). "Input specificity and the propagation of idiosyncratic shocks in production networks". In: *The Quarterly Journal of Economics* 131.3, pp. 1543–1592.

Bakaya, Yusuf Soner, Timur Hülagü, and Hande Küçük (2013). "Oil price uncertainty in a small open economy". In: *IMF Economic Review* 61.1, pp. 168–198.

Basu, Susanto and Brent Bundick (2017). "Uncertainty shocks in a model of effective demand". In: *Econometrica* 85.3, pp. 937–958.

Blaum, Joaquin, Federico Esposito, and Sebastian Heise (2025). "Input sourcing under supply chain risk: evidence from us manufacturing firms". In: *FRB of New York Staff Report* 1141.

— (2026). "Input Sourcing under Climate Risk: Evidence from US Manufacturing Firms". In: *Review of Economic Studies*, rdag031.

Blaum, Joaquin, Claire Lelarge, and Michael Peters (2018). "The gains from input trade with heterogeneous importers". In: *American Economic Journal: Macroeconomics* 10.4, pp. 77–127.

Bloom, Nicholas (2009). "The impact of uncertainty shocks". In: *econometrica* 77.3, pp. 623–685.

— (2014). "Fluctuations in uncertainty". In: *Journal of economic Perspectives* 28.2, pp. 153–176.

Bloom, Nicholas et al. (2018). "Really uncertain business cycles". In: *Econometrica* 86.3, pp. 1031–1065.

Born, Benjamin and Johannes Pfeifer (2021). "Uncertainty-driven business cycles: Assessing the markup channel". In: *Quantitative economics* 12.2, pp. 587–623.

Caporale, Guglielmo Maria, Faek Menla Ali, and Nicola Spagnolo (2015). "Oil price uncertainty and sectoral stock returns in China: A time-varying approach". In: *China Economic Review* 34, pp. 311–321.

Carvalho, Vasco M et al. (2021). "Supply chain disruptions: Evidence from the great east japan earthquake". In: *The Quarterly Journal of Economics* 136.2, pp. 1255–1321.

Castro-Vincenzi, Juanma et al. (2024). *Weathering the storm: Supply chains and climate risk*. Tech. rep. National Bureau of Economic Research.

Charoenwong, Ben, Miaozhe Han, and Jing Wu (2023). "Trade and foreign economic policy uncertainty in supply chain networks: Who comes home?" In: *Manufacturing & Service Operations Management* 25.1, pp. 126–147.

Cheng, Dong et al. (2019). "How does the Chinese economy react to uncertainty in international crude oil prices?" In: *International Review of Economics & Finance* 64, pp. 147–164.

David, Joel M, Lukas Schmid, and David Zeke (2022). "Risk-adjusted capital allocation and misallocation". In: *Journal of Financial Economics* 145.3, pp. 684–705.

De Loecker, Jan (2011). "Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity". In: *Econometrica* 79.5, pp. 1407–1451.

De Paoli, Bianca (2009). "Monetary policy and welfare in a small open economy". In: *Journal of international Economics* 77.1, pp. 11–22.

Dhyne, Emmanuel et al. (2023). "Endogenous production networks with fixed costs". In: *Journal of International Economics* 145, p. 103841.

Dixit, Avinash (1989a). "Entry and exit decisions under uncertainty". In: *Journal of political Economy* 97.3, pp. 620–638.

— (1989b). "Hysteresis, import penetration, and exchange rate pass-through". In: *The Quarterly Journal of Economics* 104.2, pp. 205–228.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu (2015). "Competition, markups, and the gains from international trade". In: *American Economic Review* 105.10, pp. 3183–3221.

Elder, John and Apostolos Serletis (2010). "Oil price uncertainty". In: *Journal of money, credit and banking* 42.6, pp. 1137–1159.

Esposito, Federico (2022). "Demand risk and diversification through international trade". In: *Journal of International Economics* 135, p. 103562.

Esposito, Federico, Sebastian Heise, and Joaquin Blaum (2024). "Input Sourcing under Climate Risk: Evidence from US Manufacturing Firms". In: *Available at SSRN 5127342*.

Ferderer, J Peter (1996). "Oil price volatility and the macroeconomy". In: *Journal of macroeconomics* 18.1, pp. 1–26.

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Keith Kuester, et al. (2015). "Fiscal volatility shocks and economic activity". In: *American Economic Review* 105.11, pp. 3352–3384.

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Juan F Rubio-Ramirez, et al. (2011). "Risk matters: The real effects of volatility shocks". In: *American Economic Review* 101.6, pp. 2530–2561.

Fornaro, Luca and Martin Wolf (2023). "The scars of supply shocks: Implications for monetary policy". In: *Journal of Monetary Economics* 140, S18–S36.

Gali, Jordi and Tommaso Monacelli (2005). "Monetary policy and exchange rate volatility in a small open economy". In: *The Review of Economic Studies* 72.3, pp. 707–734.

Gervais, Antoine (2018). "Uncertainty, risk aversion and international trade". In: *Journal of International Economics* 115, pp. 145–158.

Gorodnichenko, Yuriy and Michael Weber (2016). "Are sticky prices costly? Evidence from the stock market". In: *American Economic Review* 106.01, pp. 165–199.

Grossman, Gene M, Elhanan Helpman, and Hugo Lhuillier (2023). "Supply chain resilience: Should policy promote international diversification or reshoring?" In: *Journal of Political Economy* 131.12, pp. 3462–3496.

Grossman, Gene M, Elhanan Helpman, and Stephen J Redding (2024). "When tariffs disrupt global supply chains". In: *American Economic Review* 114.4, pp. 988–1029.

Handley, Kyle et al. (2024). "Firm input choice under trade policy uncertainty". In: *Journal of International Economics* 150, p. 103909.

Huneeus, Federico (2018). "Production network dynamics and the propagation of shocks". In: *Graduate thesis, Princeton University, Princeton, NJ* 52.

Jo, Soojin (2014). "The effects of oil price uncertainty on global real economic activity". In: *Journal of Money, Credit and Banking* 46.6, pp. 1113–1135.

Justiniano, Alejandro and Giorgio E Primiceri (2008). "The time-varying volatility of macroeconomic fluctuations". In: *American Economic Review* 98.3, pp. 604–641.

Kopytov, Alexandr et al. (2022). "Endogenous production networks under supply chain uncertainty". In: *Available at SSRN*.

— (2024). "Endogenous production networks under supply chain uncertainty". In: *Econometrica* 92.5, pp. 1621–1659.

Lee, Kiseok, Shawn Ni, and Ronald A Ratti (1995). "Oil shocks and the macroeconomy: the role of price variability". In: *The Energy Journal* 16.4, pp. 39–56.

Levchenko, A et al. (2016). *Aggregate fluctuations in adaptive production networks*. Tech. rep. Technical report, Working Paper Series, University of Zurich, Department of

Lewellen, Katharina (2006). "Financing decisions when managers are risk averse". In: *Journal of financial economics* 82.3, pp. 551–589.

Miyamoto, Wataru and Thuy Lan Nguyen (2024). "International input–output linkages and changing business cycle volatility". In: *Journal of International Economics* 147, p. 103869.

Muris, Chris et al. (2023). *Inventory, Sourcing, and the Effects of Trade Costs: Theory and Empirical Evidence*. Tech. rep. CESifo Working Paper.

Nakata, Taisuke (2017). "Uncertainty at the zero lower bound". In: *American Economic Journal: Macroeconomics* 9.3, pp. 186–221.

Nekarda, Christopher J and Valerie A Ramey (2019). *The cyclicality of the price-cost markup*. Tech. rep. Working Paper.[1026].

Plante, Michael and Nora Traum (2012). "Time-varying oil price volatility and macroeconomic aggregates". In: *Center for Applied Economics and Policy Research Working Paper* 2012-002.

Punzi, Maria Teresa (2019). "The impact of energy price uncertainty on macroeconomic variables". In: *Energy policy* 129, pp. 1306–1319.

Restuccia, Diego and Richard Rogerson (2008). "Policy distortions and aggregate productivity with heterogeneous establishments". In: *Review of Economic dynamics* 11.4, pp. 707–720.

Xu, Qinhua, Buben Fu, and Bin Wang (2022). "The effects of oil price uncertainty on Chinas economy". In: *Energy Economics* 107, p. 105840.

Zhang, Xiaoyu, Jinlan Zhou, and Xiaodong Du (2022). "Impact of oil price uncertainty shocks on Chinas macro-economy". In: *Resources Policy* 79, p. 103080.

A Appendix Regression

A.0.1 Within HS2

Table 15: Risk Shifting Behavior, Manufacturing

	(1) #Suppliers	(2) #Origin Countries	(3) #Ship	(4) #TEU
Delivery Risk	0.218***	0.212***	0.409***	0.479***
	(11.11)	(11.89)	(12.10)	(11.75)
L.Delivery Risk	0.0731***	0.0657***	0.132***	0.153***
	(5.08)	(4.64)	(4.17)	(3.90)
L2.Delivery Risk	0.0450***	0.0420**	0.0763**	0.0826*
	(3.75)	(3.10)	(3.02)	(2.34)
L3.Delivery Risk	0.0561***	0.0457***	0.0652*	0.0791
	(4.02)	(3.71)	(2.11)	(2.03)
Delivery Time	-0.0158*	-0.0186*	-0.0352*	-0.0343*
	(-2.06)	(-2.72)	(-2.40)	(-2.17)
L.Delivery Time	-0.0223**	-0.0230***	-0.0377*	-0.0346*
	(-2.93)	(-4.10)	(-2.47)	(-2.06)
L2.Delivery Time	-0.0266**	-0.0254**	-0.0392*	-0.0497*
	(-3.38)	(-3.42)	(-2.37)	(-2.64)
L3.Delivery Time	-0.0137	-0.00929	-0.0298	-0.0268
	(-1.90)	(-1.45)	(-1.97)	(-1.54)
Cumulative : Delivery Risk	0.3923***	0.3658***	0.6820***	0.7941***
SE	0.0354	0.0339	0.0644	0.0794
Cumulative : Delivery Delay	-0.0784***	-0.0762***	-0.1418***	-0.1455***
SE	0.0128	0.0125	0.0238	0.0286
N	121,436	121,436	121,436	121,436
Adj. Rš	0.720	0.740	0.733	0.729
FirmŒHS2 FE	yes	yes	yes	yes
HS2ŒTime FE	yes	yes	yes	yes

Notes: This regression is clustered at the firm and time level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 16: Risk Shifting Behavior, Manufacturing

	(1) #Suppliers	(2) #Countries	(3) #Ship	(4) TEU
Delivery Uncertainty	0.198***	0.189***	0.369***	0.434***
	(10.43)	(10.58)	(10.56)	(9.81)
Big=1 × Delivery Uncertainty	0.0393**	0.0356**	0.0708**	0.100**
	(2.38)	(2.20)	(2.37)	(2.43)
N	163,729	163,729	163,729	163,729
Adj. Rš	0.723	0.739	0.733	0.718
FirmŒHS2 FE	yes	yes	yes	yes
HS2ŒTime FE	yes	yes	yes	yes

Notes: This regression is clustered at time and firm level

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B Appendix

B.1 Appendix: Approximation of Expected Unit Cost $E(c_i)$

Because the outer production aggregator is CES, the ex-post unit cost is a nonlinear function of the material sourcing term $\Theta_{M,i}$. To obtain a tractable expression for expected unit cost, I approximate $E[c_i]$ using a second-order Taylor expansion around the mean of $\Theta_{M,i}$. The resulting expression links $E[c_i]$ to the mean and variance of $\Theta_{M,i}$, capturing how both average sourcing capability and delivery-risk dispersion affect expected production costs.

The unit cost of firm i can be written as

$$c_i = \frac{1}{\phi_i}\left[\alpha_L^{\sigma_{LM}} w^{1-\sigma_{LM}} + (1-\alpha_L)^{\sigma_{LM}} \Theta_{M,i}^{\frac{1-\sigma_{LM}}{1-\sigma_M}}\right]^{\frac{1}{1-\sigma_{LM}}}.$$

For notational convenience, define

$$A_L \equiv \alpha_L^{\sigma_{LM}} w^{1-\sigma_{LM}}, \qquad A_M \equiv (1-\alpha_L)^{\sigma_{LM}},$$

and

$$\kappa \equiv \frac{1-\sigma_{LM}}{1-\sigma_M}, \qquad \gamma \equiv \frac{1}{1-\sigma_{LM}}.$$

Then unit cost can be expressed as a function of Θ:

$$c_i(\Theta) = \frac{1}{\phi_i}\left[A_L + A_M \Theta^{\kappa}\right]^{\gamma}.$$

Let

$$\bar{\Theta}_i = E[\Theta_{M,i}], \qquad \sigma^2_{\Theta,i} = \text{Var}(\Theta_{M,i}).$$

A second-order Taylor expansion of $c_i(\Theta)$ around $\bar{\Theta}_i$ gives

$$E[c_i] \approx c_i(\bar{\Theta}_i) + \frac{1}{2} c_i''(\bar{\Theta}_i)\sigma^2_{\Theta,i}.$$

The second derivative of $c_i(\Theta)$ is

$$c_i''(\Theta) = \frac{1}{\phi_i}\gamma A_M \kappa \left[(\kappa-1)\Theta^{\kappa-2}\left(A_L + A_M\Theta^{\kappa}\right)^{\gamma-1} + (\gamma-1)A_M\kappa\Theta^{2\kappa-2}\left(A_L + A_M\Theta^{\kappa}\right)^{\gamma-2}\right].$$

Substituting this expression into the Taylor approximation yields

$$\begin{aligned} E[c_i] \approx & \frac{1}{\phi_i}\left[A_L + A_M\bar{\Theta}_i^{\kappa}\right]^{\gamma} \\ & + \frac{1}{2}\frac{1}{\phi_i}\gamma A_M\kappa\left[(\kappa-1)\bar{\Theta}_i^{\kappa-2}\left(A_L + A_M\bar{\Theta}_i^{\kappa}\right)^{\gamma-1} + (\gamma-1)A_M\kappa\bar{\Theta}_i^{2\kappa-2}\left(A_L + A_M\bar{\Theta}_i^{\kappa}\right)^{\gamma-2}\right]\sigma^2_{\Theta,i}. \end{aligned}$$

This approximation summarizes how both the mean and variance of the sourcing term $\Theta_{M,i}$ affect expected unit cost when the outer production aggregator between labor and intermediate inputs is CES.

B.2 Appendix: Approximate $E(\Theta_M)$ and $Var(\Theta_M)$

Consider a final-good firm i with chosen supplier network Z_i. In the benchmark specification, supplier weights are normalized within the chosen network. Let $a_{si} > 0$ denote the raw supplier-specific weight. For $s \in Z_i$, define

$$\alpha_{si}(Z_i) = \frac{a_{si}}{\sum_{r\in Z_i} a_{ri}}, \qquad \sum_{s\in Z_i}\alpha_{si}(Z_i) = 1.$$

The material price-index term can be written as

$$\Theta_{M_i} = \sum_{s\in Z_i}\alpha_{si}(Z_i)\, p_s^{1-\sigma_M}\Gamma_{si}^{\sigma_M-1}.$$

where $\alpha_{si}(Z_i) = \frac{a_{si}^{\sigma_M-1}}{\sum_{r\in Z_i} a_{ri}^{\sigma_M-1}}$ is the normalized supplier weights given chosen supplier network Z_i and $\sum_{s\in Z_i}\alpha_{si}(Z_i) = 1$

Define

$$A_{si}(Z_i) \equiv \alpha_{si}(Z_i) p_s^{1-\sigma_M}, \qquad X_{si} \equiv \Gamma_{si}^{\sigma_M - 1}.$$

Then

$$\Theta_{M_i} = \sum_{s \in Z_i} A_{si}(Z_i) X_{si}.$$

Let

$$A_i(Z_i) = (A_{1i}(Z_i), \ldots, A_{n_i i}(Z_i))^\top, \qquad X_i = (X_{1i}, \ldots, X_{n_i i})^\top,$$

where n_i is the number of suppliers in Z_i. Then

$$\Theta_{M_i} = A_i(Z_i)^\top X_i.$$

Expectation. Conditional on the chosen supplier network Z_i,

$$E[\Theta_{M_i} \mid Z_i] = \sum_{s \in Z_i} A_{si}(Z_i) E[X_{si}].$$

In matrix form, let

$$\mu_i = E[X_i] = (E[X_{1i}], \ldots, E[X_{n_i i}])^\top.$$

Then

$$E[\Theta_{M_i} \mid Z_i] = A_i(Z_i)^\top \mu_i.$$

Variance. Because Θ_{M_i} is linear in X_i conditional on Z_i,

$$Var(\Theta_{M_i} \mid Z_i) = \sum_{s \in Z_i} \sum_{r \in Z_i} A_{si}(Z_i) A_{ri}(Z_i) Cov(X_{si}, X_{ri}).$$

Let

$$\Sigma_i = Cov(X_i).$$

Then

$$Var(\Theta_{M_i} \mid Z_i) = A_i(Z_i)^\top \Sigma_i A_i(Z_i).$$

B.3 Appendix: The Covariance Matrix of Supplier Shocks

For a given buyer, suppress the buyer index and let the supplier shock be

$$\Gamma_s = \begin{cases} \underline{\Gamma}_s, & \text{with probability } \rho_s, \\ 1, & \text{with probability } 1 - \rho_s, \end{cases} \qquad \text{where } \underline{\Gamma}_s \in (0,1).$$

Since the material-price index depends on

$$X_s = \Gamma_s^{\sigma_M - 1},$$

define

$$\underline{X}_s \equiv \underline{\Gamma}_s^{\sigma_M - 1}.$$

Then

$$X_s = \begin{cases} \underline{X}_s, & \text{with probability } \rho_s, \\ 1, & \text{with probability } 1 - \rho_s. \end{cases}$$

Therefore,

$$E[X_s] = \rho_s \underline{X}_s + (1 - \rho_s),$$

and

$$E[X_s^2] = \rho_s \underline{X}_s^2 + (1 - \rho_s).$$

Thus,

$$Var(X_s) = E[X_s^2] - E[X_s]^2 = \rho_s(1 - \rho_s)(\underline{X}_s - 1)^2.$$

For two suppliers s and r,

$$Cov(X_s, X_r) = E[X_s X_r] - E[X_s]E[X_r].$$

Let

$$\pi_{sr} \equiv Pr(\Gamma_s = \underline{\Gamma}_s, \Gamma_r = \underline{\Gamma}_r)$$

denote the joint probability that both suppliers are hit by the shock. Then

$$E[X_s X_r] = \pi_{sr} \underline{X}_s \underline{X}_r + (\rho_s - \pi_{sr})\underline{X}_s + (\rho_r - \pi_{sr})\underline{X}_r + (1 - \rho_s - \rho_r + \pi_{sr}).$$

Equivalently,

$$Cov(X_s, X_r) = (\pi_{sr} - \rho_s \rho_r)(\underline{X}_s - 1)(\underline{X}_r - 1).$$

In the general case, the covariance matrix is

$$\Sigma = Cov(X) = \begin{bmatrix} Var(X_1) & Cov(X_1, X_2) & \cdots & Cov(X_1, X_n) \\ Cov(X_2, X_1) & Var(X_2) & \cdots & Cov(X_2, X_n) \\ \vdots & \vdots & \ddots & \vdots \\ Cov(X_n, X_1) & Cov(X_n, X_2) & \cdots & Var(X_n) \end{bmatrix}.$$

In the current quantitative implementation, supplier shocks are independent conditional on their disruption probabilities. Therefore,

$$\pi_{sr} = \rho_s \rho_r \qquad \text{for } s \neq r,$$

which implies

$$Cov(X_s, X_r) = 0 \qquad \text{for } s \neq r.$$

Thus, in the current implementation, Σ is diagonal.

B.4 Appendix: Supplier Productivity and Heterogeneous Cost Responses

The same delivery-risk shock can generate different cost responses across firms when they source from suppliers with different productivity levels. Consider the unit cost function

$$c_i(\Theta_i) = \frac{1}{\phi_i}\left[A_L + A_M \Theta_i^{\kappa}\right]^{\gamma},$$

where

$$A_L = \alpha_L^{\sigma_{LM}} w^{1-\sigma_{LM}}, \qquad A_M = (1-\alpha_L)^{\sigma_{LM}}, \qquad \kappa = \frac{1-\sigma_{LM}}{1-\sigma_M}, \qquad \gamma = \frac{1}{1-\sigma_{LM}}.$$

Ignoring the variance, the elasticity of unit cost with respect to the material sourcing term is

$$\frac{\partial \log c_i}{\partial \log \Theta_i} = \gamma\kappa \frac{A_M \Theta_i^{\kappa}}{A_L + A_M \Theta_i^{\kappa}} = \frac{1}{1-\sigma_M}\omega_{M,i},$$

where

$$\omega_{M,i} = \frac{A_M \Theta_i^{\kappa}}{A_L + A_M \Theta_i^{\kappa}}$$

is the material cost share. Since $\sigma_M > 1$ and $\sigma_{LM} < 1$ in the calibration, $\kappa < 0$. Therefore,

$$\frac{\partial \omega_{M,i}}{\partial \log \Theta_i} = \kappa \omega_{M,i}(1-\omega_{M,i}) < 0.$$

When there is no uncertainty, final good producer will only choose the most productive supplier. A firm paired with more productive supplier will have a higher value of Θ_i, corresponding to a better material sourcing position, has a lower material cost share and therefore a smaller elasticity of total unit cost with respect to changes in the material sourcing term. Thus, even when all suppliers face the same disruption probability and supplier networks are held fixed, firms with better initial sourcing positions experience smaller increases in ex-

pected marginal cost when delivery uncertainty rises. This heterogeneity in cost elasticities generates reallocation across firms and contributes to an increase in the aggregate markup.