Supply Chain Shocks and the Macroeconomy: Evidence from Global Shipping Disruptions*

Diego R. Känzig†   Ramya Raghavan‡

July, 2026

Abstract

This paper studies the macroeconomic consequences of global supply chain disruptions, focusing on maritime choke points critical to international trade. We identify supply chain shocks by combining narrative accounts of disruptions at two key locations—the Suez and the Panama Canals—with high-frequency financial data on shipping rates. Supply chain shocks lead to a persistent increase in shipping costs, causing a decline in economic activity and an increase in producer and consumer prices. Shipping capacity initially contracts before adjusting sluggishly. The shocks also lengthen delivery times and increases input shortages, while leaving geopolitical risk largely unchanged. Using granular input-output data, we document heterogeneous sectoral effects and estimate substitution elasticities, finding limited short-run but stronger medium-run substitution.

JEL classification: E30, E31, F60, R40

Keywords: Supply chains, disruptions, shipping costs, macroeconomic effects, substitution

*We thank Marios Angeletos, Florin Bilbiie, Luigi Bocola, Refet Gürkaynak, Marty Eichenbaum, Matteo Iacoviello, Ishan Nath, Marco Del Negro, Giorgio Primiceri, Hélène Rey, Kunal Sangani, Jon Steinsson, Christian Wolf, Raf Wouters, Francesco Zanetti, and participants at the EC CEPR JIE conference, the SED conference, the NBER Summer Institute, the Swiss National Bank, and the International Monetary Fund for helpful comments and suggestions.

†Northwestern University, CEPR and NBER. E-mail: dkaenzig@northwestern.edu. Web: www.diegokaenzig.com.

‡Northwestern University. E-mail: ramya.raghavan@u.northwestern.edu. Web: www.ramyaraghavan.com.

1. Introduction

Supply chains have become increasingly interconnected over the last decades. Recent events such as the Covid pandemic, Russia's invasion of Ukraine, and extreme weather have all highlighted the fragility of global supply chains. Such disruptions pose significant risks to the global economy, with potentially profound implications for growth, inflation, and employment across countries. Understanding the macroeconomic consequences of supply chain shocks is therefore crucial for designing policy responses and strengthening the resilience of global supply chains.

In this paper, we provide new evidence on the macroeconomic implications of supply chain disruptions. Our identification strategy leverages the fact that global supply chains critically rely on maritime trade, which in turn heavily depends on a small number of strategic choke points. We focus on two of the most important ones: the Suez Canal and the Panama Canal. We compile a comprehensive narrative account of disruptive incidents at these locations, such as groundings, collisions, and extreme weather events. Such events can cause major disruptions to the global shipping network and are plausibly exogenous to the global economy. We then isolate the market impact of these disruptions by measuring changes in shipping rates in narrow windows around the events, exploiting high-frequency financial data. To account for potential predictability, we orthogonalize the surprises with respect to macroeconomic and financial information available before the events. Using the resulting series as an instrument in a semi-structural model of the global economy, we identify a structural supply chain shock.

An adverse supply chain shock causes a stark and persistent increase in shipping rates, with effects that extend well beyond the initial disruption. Even short-lived disruptions can trigger ripple effects lasting for months, as ship rerouting reduces effective capacity, creates bottlenecks, and disrupts the tightly coordinated global shipping network. Higher shipping costs pass through to commodity prices, which increase significantly with some lag. Global shipping capacity initially contracts before expanding sluggishly in response to persistently elevated shipping costs. Importantly, the shock does not lead to any significant changes in geopolitical risk. This is reassuring, since we deliberately abstract from disruptions related to geopolitical tensions, which may affect the global economy through channels unrelated to supply chains.

Global supply chain shocks have meaningful effects on the U.S. economy. Industrial production falls significantly while consumer prices rise persistently. Quantitatively, a shock that raises shipping rates by 10 percent increases commodity prices by about 2 percent, lowers U.S. industrial production by around 0.5 percent, and raises the U.S. CPI by close to 0.4 percent. These stagflationary effects create a trade-off for monetary policy, reflected in the ambiguous response of the short-term interest rate. Finally, the shock leads to a significant depreciation of the dollar.

We also show that the shock leads to a significant increase in supplier delivery times and measures of industry shortages, corroborating its interpretation as a supply chain shock. The sluggish response of commodity prices, together with the delayed adjustment in shipping capacity, further suggests that we are not merely picking up commodity price shocks.

A comprehensive series of sensitivity checks shows that the results are robust across several dimensions, including the model specification, the sample period, and the identification strategy. In particular, the responses are similar when we relax the invertibility requirement and when we estimate the dynamic effects using local projections.

Using the identified supply chain shocks, we investigate the role of supply chain pressures in the recent inflationary episode. We find that these shocks contributed meaningfully to inflation dynamics over the 2020–2022 period. They explain a considerable share of the rise in inflation through 2021, consistent with widespread logistical bottlenecks and maritime stress during that period. However, they do not account for the bulk of the inflation peak observed in 2022. This suggests that supply chain shocks were important early in the episode, while other forces—including fiscal stimulus, accommodative monetary policy, and energy price shocks—played a more prominent role later on.

Counterfactual analyses suggest that monetary policy plays an important role in the transmission of supply chain shocks. A more aggressive monetary response could largely stabilize prices in the face of such shocks, but at the cost of a significantly steeper decline in output.

Moving beyond aggregate impacts, we document substantial heterogeneity in the transmission of supply chain shocks across sectors. Using disaggregated data on industrial production and producer prices, we show that the effects vary systematically with sectors' exposure to intermediate inputs and global supply chains. Price pressures are particularly strong in upstream industries such as petroleum and metals. Output contractions, by contrast, are concentrated in downstream sectors, most notably motor vehicles, and in input-intensive non-durable industries such as chemicals and plastics.

Motivated by these patterns, we estimate macro elasticities of substitution across intermediate inputs using granular input-output data. Our approach relates the shock-induced responses of expenditure shares across supplier-purchaser pairs to the corresponding responses of relative input prices. This allows us to recover horizon-specific elasticities that capture equilibrium reallocation in response to supply chain shocks, distinguishing between short-run and medium-run adjustment. We find that substitution is limited in the short run but rises above one after one year. These estimates point to limited immediate substitution, but economically meaningful reallocation of sourcing across suppliers over the medium run.

Related literature and contribution. This paper contributes to a rapidly growing literature studying the economic effects of supply chain disruptions, from both empirical and theoretical perspectives.

A recent empirical literature studies the effects of supply chain pressures and shipping cost surges on the macroeconomy (Herriford et al., 2016; Jacks and Stuermer, 2021; Carrière-Swallow et al., 2023, among others). Several studies construct direct measures of supply chain pressures, using real-time congestion of containerships in major ports (Bai et al., 2024), route-specific shipping disruptions (Bai et al., 2025), or by combining purchasing managers' indices with transportation costs (Benigno et al., 2022), to assess how supply chain shocks affect economic outcomes. Caldara, Iacoviello, and Yu (2024) construct a newspaper-based measure of input shortages and find that such shortages are associated with persistent inflationary effects. The closest paper to ours is Bini (2025), who identifies global supply chain shocks using narrative information from price surcharge announcements by major container shipping companies. Our approach is complementary but differs along two key dimensions. First, focusing on disruptive events at major maritime choke points allows us to study a longer sample, increasing statistical power. Second, using high-frequency movements in shipping rates around these events helps us better account for confounding shocks.

At a more disaggregated level, Blaum, Esposito, and Heise (2026) construct a measure of supply chain risk at the firm level based on transaction data on U.S. manufacturing imports and show that unanticipated shipping delays are associated with lower firm-level employment and profits. Castro-Vincenzi et al. (2024) study how firms structure supply chains under climate risk. We contribute to these literatures by providing new evidence on the aggregate and sectoral effects of supply chain disruptions, and by estimating how sourcing patterns adjust in response to these shocks.

Methodologically, our approach builds on the literature on high-frequency identification, which was developed in the monetary policy setting (Kuttner, 2001; Gürkaynak, Sack, and Swanson, 2005; Gertler and Karadi, 2015; Nakamura and Steinsson, 2018, among others) and more recently employed in the context of oil and carbon markets (Känzig, 2021; Känzig, 2023). In this literature, policy surprises are identified using high-frequency asset price movements around policy events, such as FOMC or OPEC meetings. The idea is to isolate the impact of policy news by measuring the change in asset prices in a tight window around the events. We employ this approach in the context of global supply chains, allowing for credible identification under relatively weak structural assumptions.

Finally, our paper relates to a growing literature estimating substitution elasticities in production. Existing estimates generally point to limited short-run substitutability across material inputs or suppliers, both at the industry level (Atalay, 2017; Miranda-Pinto, 2021) and at the firm level (Barrot and Sauvagnat, 2016; Boehm, Flaaen, and Pandalai-Nayar, 2019; Carvalho et al., 2021). Recent work by Peter and Ruane (2025) shows that intermediate-input substitutability can be substantially larger over longer horizons and in response to more persistent shocks. Evidence from trade also points to horizon as an important dimension: Boehm, Levchenko, and Pandalai-Nayar (2023) estimate substantially larger long-run than short-run trade elasticities, with adjustment unfolding over several years. Related dynamic network models emphasize that real and nominal adjustment frictions can slow the response of production networks to shocks (Huneeus, 2018; Liu and Tsyvinski, 2020; Minton and Wheaton, 2023). We contribute by estimating horizon-specific, reduced-form macro elasticities of substitution across intermediate inputs in response to identified supply chain shocks. These estimates capture equilibrium reallocation across suppliers and sectors and can help discipline dynamic production-network models.

In this sense, we also contribute to an influential literature that emphasizes production networks as a propagation channel for microeconomic and sectoral shocks (Acemoglu, Akcigit, and Kerr, 2016; Baqaee and Farhi, 2019; Carvalho and Tahbaz-Salehi, 2019, among others). Related empirical work documents how disruptions propagate through firm and sectoral input linkages (Barrot and Sauvagnat, 2016; Boehm, Flaaen, and Pandalai-Nayar, 2019; Carvalho et al., 2021). Recent contributions have shown how supply chain disruptions and granular bottlenecks can generate large aggregate effects (Di Giovanni et al., 2022; Comin, Johnson, and Jones, 2023; Alessandria et al., 2023; Acemoglu and Tahbaz-Salehi, 2024). Consistent with this view, we find that disruptions to global shipping choke points have sizeable macroeconomic and sectoral effects.

Outline. The paper proceeds as follows. Section 2 describes our strategy to identify supply chain shocks. We discuss the role of maritime choke points in global supply chains, provide a detailed narrative account of disruptions at these choke points, and construct shipping cost surprises using high-frequency data on shipping rates. Section 3 introduces our empirical framework. Section 4 documents the aggregate macroeconomic effects of supply chain shocks. Section 5 studies their sectoral impacts and estimates macro elasticities of substitution across intermediate inputs using granular data. Section 6 concludes.

2. Identification Strategy

Global supply chains are heavily reliant on maritime trade, with over 80 percent of the world's trade by volume transported by sea. As a consequence, shipping rates offer a timely barometer of supply chain conditions.

Figure 1 shows the evolution of shipping rates, as proxied by the Baltic Dry Index (BDI), from the 1970s to date. The BDI tracks the cost of transporting bulk commodities like iron ore, coal, and grain across major shipping routes and is commonly used as a proxy for global shipping rates.

Figure 1: Global Shipping Rates as a Barometer of Supply Chain Pressures

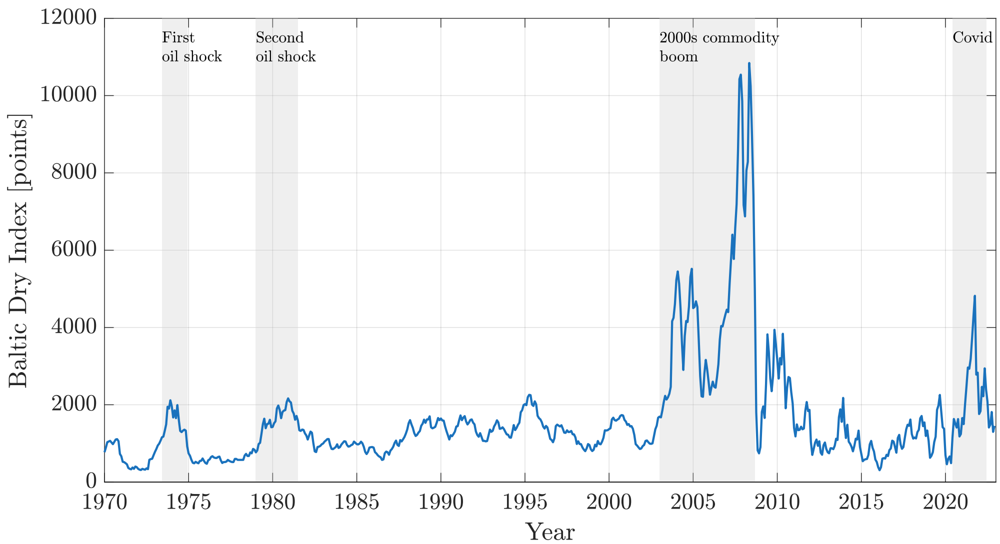

Notes: Global shipping costs, proxied by the Baltic Dry Index published by the Baltic Exchange. The index tracks the cost of transporting bulk commodities across major shipping routes.

Fluctuations in shipping rates are driven by changes in global demand and supply. The sharp spikes during the oil shocks of the 1970s were clearly driven by supply disruptions. The surge in the 2000s, by contrast, can be largely attributed to booming Chinese demand for commodities. The spike around the Covid-19 period likely reflected a mix of both—disrupted logistics and a rapid rebound in global goods demand. The fact that shipping rates are driven by both demand and supply pressures complicates the identification of supply chain shocks based on shipping price movements. Simply regressing economic outcomes on shipping costs will in general lead to biased results.

We address this challenge by leveraging the reliance of global supply chains on critical choke points combined with high-frequency financial data to isolate some variation in shipping rates that is driven by supply chain disruptions and plausibly exogenous to the world economy.

2.1. Choke Points in Global Supply Chains

Maritime trade facilitates the movement of goods across vast distances at a relatively low cost and is crucial to the functioning of international commerce. Figure 2 displays the main global sea trade routes together with associated shipping times. A key feature of these trade routes is the stark dependence on certain critical points, known as maritime choke points.

Maritime choke points are narrow waterways that connect global sea trade routes. Such choke points are characterized by a narrow width, high shipping traffic, and few viable alternate routes. Disruptions at these choke points have far-reaching consequences for global supply chains. The two most critical choke points are the Suez Canal and the Panama Canal, which are the key focus of this paper. We start by providing some background information on these two choke points.

Suez Canal. The Suez Canal, located in Egypt, is a maritime choke point that connects the Mediterranean Sea and Red Sea through the Gulf of Suez. The Suez Canal provides the shortest sea trade route between Europe and Asia. For instance, vessels traveling through the Suez Canal for journeys between the Persian Gulf and Amsterdam-Rotterdam-Antwerp save 15 days of transit time in comparison to alternative routes (see Figure 2). Over 15 percent of international maritime trade passes through the Suez Canal, making international trade highly vulnerable to disruptions there.

Throughout its history, the Suez Canal has experienced numerous disruptions to shipping traffic due to its narrow, constrained passage—from vessel groundings and collisions to fires, localized maritime-security incidents, and adverse weather conditions (see Figure 3a). One such notable and recent incident occurred on March 23, 2021, when Ever Given, a container ship, ran aground and became lodged sideways across the canal for six days, blocking the trade route. However, this is only one of many disruptions that have affected this critical maritime choke point over the years. As another example, we provide below an excerpt of a Reuters news article discussing a grounding in the Suez Canal on November 8, 2004 (Reuters, 2004):

Egypt's Suez Canal has been blocked by a broken-down oil tanker and could stay shut

Figure 2: Major Global Sea Trade Routes, Choke Points, and Shipping Times

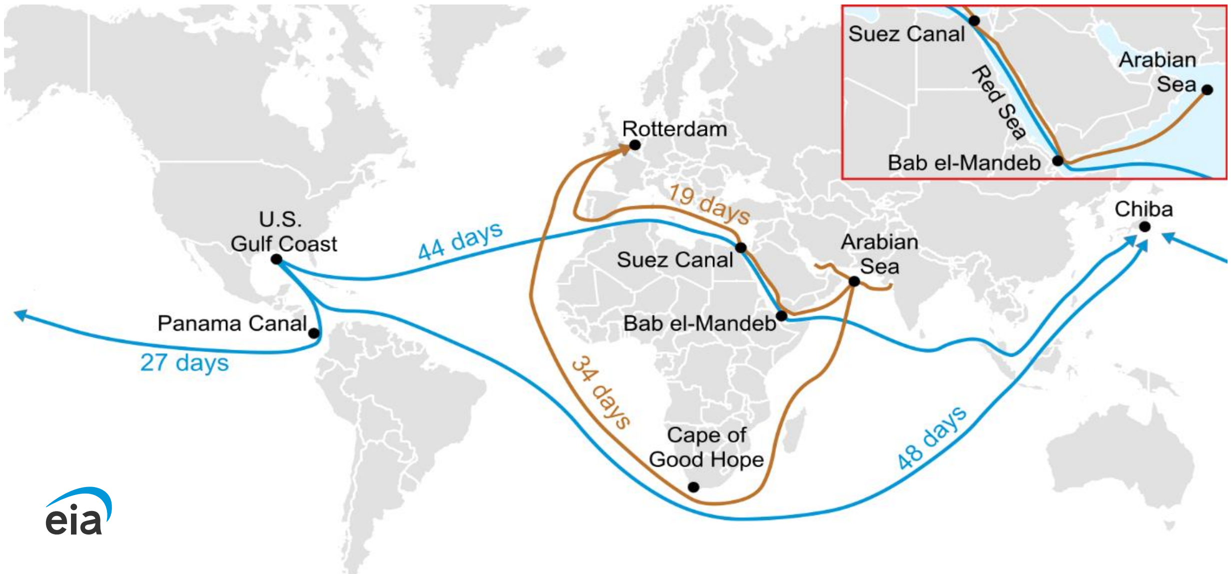

Notes: Major global sea trade routes. In brown: Vessels traveling through the Suez Canal for journeys between the Persian Gulf and Amsterdam-Rotterdam-Antwerp save 15 days of transit time in comparison to alternative routes. In blue: Vessels traveling through the Panama Canal for journeys between the U.S. Gulf Coast and Chiba, Japan save 17-21 days of transit time in comparison to alternative routes. Source: U.S. Energy Information Administration (2024).

> *for another two days (...)*
>
> *Navigation came to a standstill late on Saturday when the 154,000 deadweight-tonne Liberian-flagged vessel Tropic Brilliance, carrying a cargo of crude, ran aground while passing through the canal. (...)*
>
> *Shipping sources expected traffic to be disrupted until Wednesday at least.*

Disruptions in the Suez Canal, such as the examples above, are widely discussed across both mainstream and maritime-specific news agencies.

Panama Canal. The Panama Canal, located in Panama, is a maritime choke point that connects the Atlantic Ocean and Pacific Ocean. The canal is the shortest sea trade route between the Atlantic and Pacific Oceans and accounts for around 46 percent of the trade between Northeast Asia and the U.S. East Coast. Vessels traveling through the Panama Canal for journeys between the U.S. Gulf Coast and Chiba, Japan save 17-21 days of transit time in comparison to alternative routes (see Figure 2). More than 5 percent of global maritime trade passes through the Panama Canal.

Due to its narrow passage and reliance on freshwater for lock operations (see Figure 3b), the Panama Canal is prone to disruptions. Over the last decades, there were nu-

Figure 3: Maritime Choke Points

(a) Suez Canal (b) Panama Canal

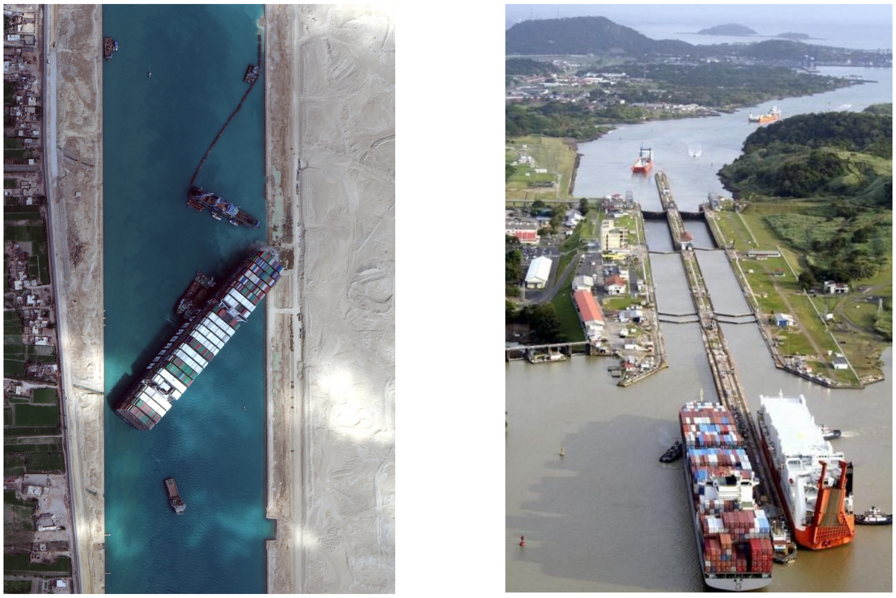

Notes: Two of the world's most critical maritime choke points. Panel (a) shows the Ever Given container ship obstructing the Suez Canal during the March 2021 grounding, which disrupted global trade for nearly a week. Panel (b) shows vessels transiting the Panama Canal's lock system, which lifts and lowers ships through a freshwater channel, via Gatun Lake, to connect the Atlantic and Pacific Oceans.

merous disruptions to shipping traffic, from adverse weather conditions and fires to vessel groundings and collisions. Recently, the shipping traffic through the Panama Canal has often been subject to transit and draft restrictions that limit the number and capacity of vessels, respectively, due to adverse weather conditions. The below excerpt from a Reuters news article discusses such a draft restriction from August 7, 2015 (Reuters, 2015):

> *The Panama Canal Authority will temporarily lower the maximum draft of ships passing through the canal, due to droughts caused by El Nino, authorities said on Friday.*
>
> *Starting on Sept. 8, the greatest draft allowed will be 39 feet (11.89 m), down from the current maximum of 39.5 feet (12.04 m), the Panama Canal Authority said.*
>
> *The change could affect about 20 percent of ships that use this route (...)*

Restrictions, as in the example above, are imposed by the Panama Canal Authority (ACP)—the authority responsible for the administration, operation, and maintenance of the canal. When restrictions are imposed, the authority issues an "Advisory to Shipping" on their official website. Such restrictions and other shipping disruptions are widely reported by both mainstream and maritime-specific news agencies.

2.2. A Narrative Analysis of Supply Chain Disruptions

We compile a comprehensive narrative dataset of events that disrupted shipping traffic through the Suez and Panama Canals using three complementary sources: specialized maritime intelligence, leading newswires, and official shipping advisories. We use Lloyd's List Intelligence to identify vessel-level casualties, such as groundings, collisions, and machinery failures. We complement and verify these records with news reports from Factiva, which capture broader disruptions, including weather-related incidents. For the Panama Canal, we further use official shipping advisories to identify operational restrictions, such as transit and draft limits.

Table 1: Disruptive Events at Maritime Choke Points

Panama Canal			Suez Canal		
Event Type	Number		Event Type	Number	
	Factiva	Lloyd's		Factiva	Lloyd's
Grounding	1	6	Grounding	62	36
Collision	1	17	Collision	6	21
Fire	1	1	Fire	5	11
Machinery Failure	0	15	Machinery Failure	0	32
Foundered/Hull Damage	0	6	Foundered/Hull Damage	0	12
El Nino/Rainfall	31	0	Weather	9	0
Landslide/Flooding	3	0	Sandstorm	8	0
Drought	5	0	Piracy/Local security	2	0
Other	3	2	Other	6	1
Total	45	47	Total	98	113

Notes: Overview of events that disrupted traffic in the Panama and the Suez Canal. The events are grouped into different categories and listed separately by source.

Our dataset identifies 303 events that disrupted shipping traffic in the two canals between 1970 and 2022, of which 211 events occurred in the Suez Canal and 92 events affected the Panama Canal. Table 1 provides an overview of the events in our dataset. The events in the Suez Canal include 98 vessel groundings, 27 collisions, and 32 machinery failures. The events in the Panama Canal include 18 collisions, 15 machinery failures, and 39 weather-related disruptions, associated with El Nino events, landslides, flooding, and drought conditions.

In selecting events, we put great care into ensuring that they are plausibly exogenous to global macroeconomic conditions. Specifically, we include vessel groundings, collisions, fires, adverse weather conditions, and piracy incidents when they are localized maritime-security events. We exclude events tied to broader geopolitical conflict, especially when their likely macroeconomic effects may operate through geopolitical-risk or oil-market channels rather than through shipping disruptions alone. Some disruptions—such as weather-related restrictions—may be partially forecastable. In these cases, isolating the unanticipated component is crucial to avoid bias from anticipatory effects.

2.3. High-frequency Identification

The identified events vary in severity and may, in some cases, be at least partly anticipated. To gauge the importance of each event while accounting for potential anticipation effects, we adopt a high-frequency identification strategy that leverages financial market data on shipping rates.

The idea is as follows. Disruptive events along global shipping routes are closely monitored by market participants and can lead to significant market reactions. Thus, we can isolate the impact of a disruptive event by measuring the change in shipping rates in a tight window around the disruption. To the extent that the global economic outlook is priced at the time of the event and unlikely to change over the short event window, we isolate some unexpected variation in shipping rates that is plausibly exogenous. Doing this systematically across all our events, we construct a series of high-frequency shipping cost surprises, i.e. the unexpected component of shipping rates associated with a disruption in shipping markets. These surprises can be used to identify a structural global supply chain shock.

Measuring shipping costs. We use the Baltic Dry Index (BDI), published daily by the London-based Baltic Exchange, to measure global shipping costs. The BDI is a composite of timecharter rates across three vessel classes—Capesize (40 percent), Panamax (30 percent), and Supramax (30 percent)—each of which reflects average shipping rates across key global dry bulk trade routes. As such, it provides a comprehensive measure of maritime transport costs for bulk commodities.

We focus on the BDI for two main reasons. First, it is a well-established barometer of global trade activity and is closely followed by market participants. Second, it is available at a daily frequency and spans a long historical sample period, making it well-suited for our time-series, high-frequency identification approach.

A potential limitation is that the BDI covers dry bulk shipping and not directly containerized freight. Data on container rates are only available at lower frequency and with more limited historical coverage. However, since freight rates tend to co-move across shipping segments due to shared cost drivers and capacity constraints, the BDI remains a broadly informative proxy for global shipping conditions.

Figure 4: Shipping Cost Surprise Series

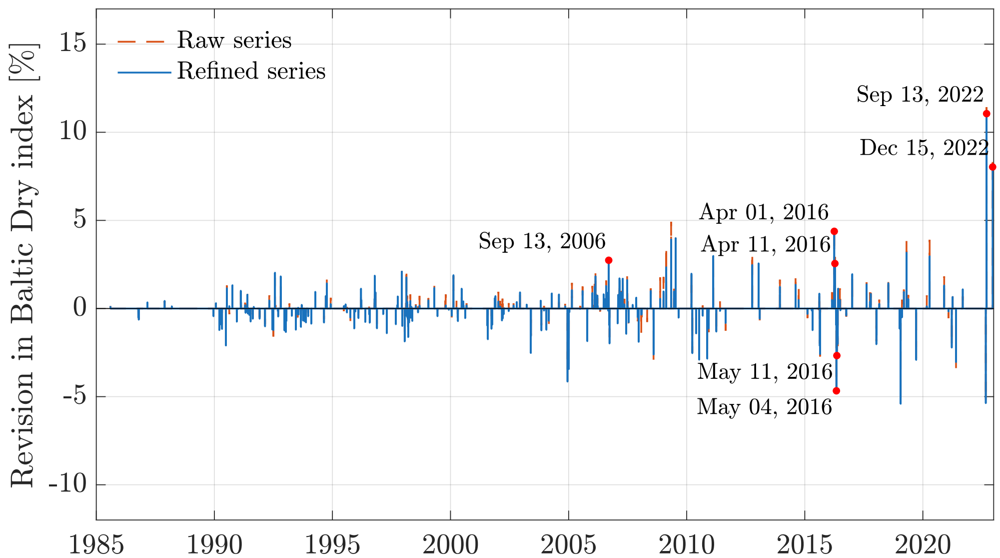

Notes: The shipping cost surprise series, computed as the daily change in the BDI around disruptive events in the Panama and Suez canals. Dashed line: raw surprise series. Solid line: refined surprise series. Dots: examples of disruptive events.

Construction of shipping cost surprises. We construct a series of shipping cost surprises by measuring how shipping costs change around the disruptive events we identify. Specifically, we record the percentage change in shipping costs on the event day compared to the last trading day before the event:[1]

$$SCSurprise_d = \frac{P_d^{\text{BDI}} - P_{d-1}^{BDI}}{P_{d-1}^{\text{BDI}}}, \tag{1}$$

where d indicates the date of the event, and P_d^{BDI} is the BDI price. Note that if an event is partially anticipated, this will isolate the unexpected component of the given event, provided that these expectations are priced in the market.

Figure 4 shows the shipping cost surprise series at the daily frequency. Events that disrupt shipping traffic can have a significant impact on shipping costs, with several events moving shipping costs in excess of 3 percent. To illustrate our approach, we discuss a few representative events. On September 13, 2006, an Egyptian dredger sank in the Suez Canal, causing a temporary closure of the waterway as a rescue operation was under way to find missing crew members. This resulted in an increase in shipping costs of 2.9

[1]If trading did not take place on the day of the event, we compute the surprise using the next trading day. Appendix C.1 shows that our results are robust to the choice of the event window.

percent.

On April 1 and April 11, 2016, the ACP extended draft restrictions in response to El Nino–related drought conditions and projected low Gatun Lake levels, causing shipping costs to rise by 4.8 percent and 2.9 percent, respectively. Subsequent rainfall improved the water outlook and allowed the authorities to postpone the restrictions twice: first on May 4, 2016, which reduced shipping costs by 4.5 percent, and again on May 11, 2016, which reduced them by 2.6 percent. The latter are examples of negative shipping cost surprises. Negative surprises may arise either from events that ease previous restrictions or from news indicating that a disruption is less severe than previously anticipated. However, focusing only on positive surprises produces very similar results (see Appendix C.1).

Finally, on September 13, 2022, a rare overflow at the Panama Canal's Gatun Locks, that temporarily blocked the west lane, resulted in an increase in shipping costs of 11.4 percent. On December 15, 2022, traffic at the Panama Canal's Miraflores locks was temporarily suspended due to a fire, again leading to an increase in shipping costs of 8.7 percent.

Predictability of shipping cost surprises. An influential literature shows that, in the monetary policy context, high-frequency surprises are predictable based on publicly available macroeconomic and financial data preceding the policy announcement (Cieslak, 2018; Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023). This predictability challenges the interpretation of such surprises as primitive "shocks" and may bias estimates of their dynamic effects.

Are shipping cost surprises also predictable based on past macroeconomic and financial variables? To assess this, we regress the daily surprise series on relevant information available before the event:

$$SCSurprise_d = \alpha + \beta' X_{d-} + \eta_d, \tag{2}$$

where d indexes days with shipping disruptions, $SCSurprise_d$ denotes the shipping cost surprise series, and X_{d-} is a set of predictors known before the event day d, as indicated by the subscript $d-$.

As predictors, we consider a wide range of macroeconomic and financial variables. For macroeconomic and financial indicators, we include the surprise components of the latest U.S. industrial production, ISM Purchasing Managers' Index (PMI), producer price index, and the trade balance releases prior to each event. These surprises are defined as the difference between the actual release and Bloomberg survey expectations. We also include the log change in the S&P 500 index and the yield curve slope, measured over the three months leading up to the event. For commodity markets, we add the three-month log changes in the WTI crude and the coal price. Finally, to account for geopolitical risk, we include the three-month change in the index by Caldara and Iacoviello (2018).

Table 2: Predictability of Shipping Cost Surprises

Shipping cost surprise:	(a) Macro news	(b) Financials	(c) Commodities	(d) Other
IP surprise	-0.184	-0.223	-0.209	-0.231
	(0.284)	(0.276)	(0.250)	(0.245)
ISM surprise	0.087	0.095	0.116	0.115
	(0.044)	(0.047)	(0.047)	(0.047)
PPI surprise	-0.000	-0.001	0.095	0.081
	(0.329)	(0.329)	(0.341)	(0.346)
Trade balance surprise	0.037	0.048	0.047	0.044
	(0.035)	(0.038)	(0.038)	(0.039)
S&P 500 (3M log change)		-0.456	0.191	0.012
		(1.454)	(1.267)	(1.266)
Yield curve slope (3M change)		-0.278	-0.205	-0.189
		(0.351)	(0.358)	(0.361)
WTI price (3M log change)			-0.209	-0.248
			(0.607)	(0.608)
Coal price (3M log change)			-1.597	-1.594
			(1.140)	(1.145)
Geopolitical risk (3M log change)				-0.001
				(0.001)
R^2	0.018	0.025	0.040	0.043
Adj. R^2	0.004	0.003	0.012	0.011

Notes: Estimated coefficients β, R^2, and adj. R^2 from predictive regressions (2) of shipping cost surprises. The predictors X are observed prior to the event and include: the surprise component of the most recent U.S. industrial production, ISM PMI, producer price index, and trade balance releases in column (a); column (b) adds the three-month change in the (log) S&P 500 and the yield curve slope; column (c) adds the three-month log change in WTI and coal price; column (d) adds three-month log change in geopolitical risk. Robust standard errors in parentheses.

Table 2 presents the results. There is limited evidence that shipping cost surprises are predictable by macroeconomic and financial variables. The R^2 is generally low, ranging from 0.02 to 0.04 across specifications. The only predictor that is statistically significant at conventional levels is the ISM surprise. This is not necessarily problematic, as it may stem from plausibly exogenous factors such as unusual weather patterns. Alternatively, a high ISM may indicate tighter supply chains, thereby making disruptions more binding or increasing their salience.

To account for this potential predictability, we follow the approach by Bauer and Swanson (2023). Specifically, we construct a refined shipping cost surprise series as the residual from the predictive regression (2), controlling for the full information set (d). The resulting series $\widetilde{SCSurprise}_d = \eta_d$, shown as the blue bars in Figure 4, closely tracks the raw series, with a correlation coefficient of close to 1.

Aggregation and additional diagnostics. Because our outcome variables of interest are only available at the monthly frequency, we aggregate the daily surprises to a monthly series as follows:

$$SCSurprise_t = \sum_{d=1}^{T_t} \frac{T_t - d + 1}{T_t} SCSurprise_{t,d}, \tag{3}$$

where T_t is the number of days in the month t. This assigns a larger weight to events occurring earlier in the month, reflecting the larger share of the month over which they can affect monthly outcomes. By construction, months without any events take value zero. We choose this weighting scheme to reflect the fact that shipping rates enter the empirical specification as daily averages over the month. Appendix C.1 shows that the results are robust to alternative aggregation schemes.

We perform a number of diagnostic checks on the surprise series as proposed in Ramey (2016), in particular with regard to autocorrelation, forecastability, and correlation with other structural shocks. We find that the surprise series is weakly autocorrelated, likely reflecting the clustering of certain types of events. Our results are unchanged when orthogonalizing the surprise series with respect to its own lags. We also find no evidence that macroeconomic or financial variables have any power in forecasting the surprise series. For all variables considered, the p-values for the Granger causality test are safely above conventional significance levels. Finally, we show that the surprise series is uncorrelated with other structural shock measures from the literature, including oil, productivity, news, monetary policy, uncertainty, financial, and fiscal policy shocks. Overall, this evidence supports the validity of the shipping cost surprise series. The corresponding figures and tables can be found in online Appendix B.

3. Econometric Approach

As shown above, the shipping cost surprise series has many desirable properties. Nonetheless, it is only an imperfect measure of the shock of interest because it does not capture all relevant disruptions in global supply chains and could be measured with error (Stock and Watson, 2018). Therefore, we do not use it as a direct shock measure but as an *instrument*. Provided that the surprise series is correlated with the supply chain shock but uncorrelated with all other shocks, we can use it to estimate the dynamic causal effects of a structural supply chain shock.

A challenge in estimating the dynamic causal effects using high-frequency surprises is the so-called power problem (Nakamura and Steinsson, 2018). Over the impulse horizon, macroeconomic variables are influenced by a myriad of other shocks, while high-frequency shipping cost surprises explain only a small share of the fluctuations in shipping rates—resulting in a low signal-to-noise ratio. This makes it difficult to directly estimate macroeconomic effects of high-frequency shipping cost surprises using local projections à la Jordà (2005).

To address this challenge, we rely on VAR techniques for estimation, using the external instruments approach (Stock and Watson, 2012; Mertens and Ravn, 2013).

3.1. Framework

We are interested in modeling global shipping markets and the U.S. economy jointly. Let $\mathbf{y}_t$ denote a $n \times 1$ vector of monthly time series. We assume that the dynamics of $\mathbf{y}_t$ can be characterized by the following structural vector moving-average representation:

$$\mathbf{y}_t = \mathbf{B}(L)\mathbf{S}\boldsymbol{\varepsilon}_t, \tag{4}$$

where $\boldsymbol{\varepsilon}_t$ is a vector of mutually uncorrelated structural shocks driving the economy, $\mathbf{B}(L) \equiv \mathbf{I} + \mathbf{B}_1 L + \mathbf{B}_2 L^2 + \ldots$ is a matrix lag polynomial, and $\mathbf{S}$ is the structural impact matrix.

Assuming that the vector-moving average process (4) is invertible, it admits the following VAR representation:

$$\mathbf{A}(L)\mathbf{y}_t = \mathbf{S}\boldsymbol{\varepsilon}_t = \mathbf{u}_t, \tag{5}$$

where $\mathbf{u}_t$ is a $n \times 1$ vector of reduced-form innovations with variance-covariance matrix $\mathrm{Var}(\mathbf{u}_t) = \boldsymbol{\Sigma}$ and $\mathbf{A}(L) \equiv \mathbf{I} - \mathbf{A}_1 L - \ldots$ is a matrix lag polynomial. Truncating the VAR to order p, we can estimate the model using standard techniques and recover an estimate of $\mathbf{A}(L)$.

We want to identify the causal impact of a single shock. Without loss of generality, let us denote the supply chain shock as the first shock in the VAR, $\varepsilon_{1,t}$. Our aim is to identify the structural impact vector $\mathbf{s}_1$, which corresponds to the first column of $\mathbf{S}$.

External instrument approach. Identification using external instruments works as follows. Suppose there is an external instrument available, z_t. In the application at hand, z_t is the shipping cost surprise series. For z_t to be a valid instrument, we need:

$$\mathbb{E}[z_t \varepsilon_{1,t}] = \alpha \neq 0 \tag{6}$$

$$\mathbb{E}[z_t \boldsymbol{\varepsilon}_{2:n,t}] = \mathbf{0}, \tag{7}$$

where $\varepsilon_{1,t}$ is the supply chain shock and $\boldsymbol{\varepsilon}_{2:n,t}$ is a $(n-1) \times 1$ vector consisting of the other structural shocks. Assumption (6) is the relevance requirement and assumption (7) is the exogeneity condition. These assumptions identify the structural impact vector $\mathbf{s}_1$ up to sign and scale:

$$\mathbf{s}_1 \equiv \boldsymbol{\theta}_0 \propto \frac{\mathbb{E}[z_t \mathbf{u}_t]}{\mathbb{E}[z_t \mathrm{u}_{1,t}]}, \tag{8}$$

provided that $E[z_t \mathrm{u}_{1,t}] \neq 0$. Note that to identify the relative impact effects, we do not need to assume invertibility. Invertibility is only required for the dynamic effects beyond impact.[2]

To facilitate interpretation, we scale the structural impact vector such that the shock corresponds to a 10 percent increase in shipping costs. We implement the estimator with a 2SLS procedure and estimate the coefficients above by regressing $\hat{\mathbf{u}}_t$ on $\hat{\mathrm{u}}_{1,t}$ using z_t as the instrument. To conduct inference, we employ a residual-based moving block bootstrap, as proposed by Jentsch and Lunsford (2019).

Relaxing VAR assumptions. The VAR approach improves precision, allowing for sharper inference. However, it relies on two potentially restrictive assumptions. The first is invertibility, meaning that the model incorporates all relevant information needed to recover the structural shocks of interest. The second pertains to the dynamic VAR structure, with the key assumption being that a finite-order VAR adequately approximates the dynamics of the data generating process. We perform two exercises to assess how restrictive these assumptions are.

First, we alternatively estimate the responses using the internal instrument approach (Plagborg-Møller and Wolf, 2019). This approach does not rely on invertibility but instead

[2]To be more precise, the VAR does not have to be fully invertible for identification with external instruments: it suffices if the shock of interest is invertible in combination with a limited lead-lag exogeneity condition (Miranda-Agrippino and Ricco, 2018).

assumes that the instrument is orthogonal to leads and lags of the structural shocks:

$$\mathbb{E}[z_t \boldsymbol{\varepsilon}_{t+j}] = \mathbf{0}, \quad \text{for } j \neq 0. \tag{9}$$

Together with the exogeneity and relevance requirement, this identifies the dynamic causal effects of interest. The approach can be implemented by augmenting the VAR with the instrument ordered first, $\bar{\mathbf{y}}_t = (z_t,\ \mathbf{y}_t')'$, and computing the impulse responses to the first orthogonalized innovation, $\bar{\mathbf{s}}_1 = [\text{chol}(\bar{\boldsymbol{\Sigma}})]_{\cdot,1}/[\text{chol}(\bar{\boldsymbol{\Sigma}})]_{1,1}$.

Second, we estimate the impulse responses to the supply chain shock using local projections. This approach directly estimates the dynamic responses and does not rely on the VAR structure, making it less prone to lag truncation bias. Specifically, we first extract the supply chain shock from the monthly VAR as $\varepsilon_{1,t} = \mathbf{s}_1' \boldsymbol{\Sigma}^{-1} \mathbf{u}_t$ (as in Stock and Watson, 2018). Next, we estimate the effects for the additional variables of interest y_i using simple local projections:

$$y_{i,t+h} = \alpha^i_{h,0} + \theta^i_h \varepsilon_{1,t} + \beta^i_{h,1} y_{i,t-1} + \ldots + \beta^i_{h,p} y_{i,t-p} + \xi_{i,t,h}, \tag{10}$$

where θ^i_h is the effect on variable i at horizon h. Note that while this is a more robust way to estimate the dynamic responses, it still retains the assumption of (partial) invertibility.

Estimating effects on additional outcome variables. To analyze the effects on a wider set of outcome variables, we adopt the following approach. For monthly variables, we augment our baseline VAR by the variable of interest and map out the response as in Gertler and Karadi (2015). For variables at lower frequencies, we rely on the local projections approach outlined above. To fix ideas, we aggregate the shock $\varepsilon_{1,t}$ by summing over the respective months before running the local projections at the lower frequency. This mitigates the problem that high-frequency instruments, when aggregated to the quarterly or annual frequency, often lack the power to credibly estimate the effects of interest (Nakamura and Steinsson, 2018).

Empirical specification. The baseline specification includes eight variables, organized into two blocks. The global supply chain block consists of real shipping cost, measured using the BDI, a real commodity price index, global shipping capacity, and a geopolitical risk index. The BDI and commodity price index are deflated by U.S. CPI. The U.S. block consists of industrial production, the consumer price index (CPI), the three-month

Treasury yield, and the real effective exchange rate.[3] Detailed information on the data, including sources, is provided in Appendix A.

We use monthly data spanning the period from 1970 to 2022 and estimate the VAR in levels, following Sims, Stock, and Watson (1990). Besides the short-term interest rate, all variables enter in log-levels. The lag order is set to 12 and the deterministic terms include a constant and a linear trend.

4. The Macroeconomic Impact of Supply Chain Shocks

How do supply chain shocks affect the global shipping market and the macroeconomy? In this section, we discuss the results based on our baseline external instruments VAR model.

First stage. The key identifying assumption in the external-instruments approach is that the instrument is relevant for the structural shock of interest and orthogonal to all other structural shocks. While this exclusion restriction is not directly testable, we carefully select the disruptions underlying our surprise series to isolate variation in shipping rates that is plausibly exogenous to macroeconomic developments.

Even if the instrument is exogenous, weak relevance can make standard inference unreliable. We therefore assess instrument strength using the first-stage regression of the shipping-cost residual from the VAR on the instrument (Montiel Olea, Stock, and Watson, 2016). The surprise series is strong: the heteroskedasticity-robust first-stage F-statistic is 33.41, well above conventional critical values.

Macroeconomic effects. Figure 5 shows the impulse responses to the identified supply chain shock, normalized to increase the real shipping cost by 10 percent on impact. In each panel, the solid black line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

A disruptive supply chain shock leads to a strong increase in shipping rates that persists for months beyond the initial event. Interestingly, these effects last significantly longer than the duration of the underlying shipping disruptions in our instrument. The main reason is that a disruptive event, even if only relatively short-lived, creates substantial ripple effects throughout the global shipping market. Ships must be rerouted,

[3] Global shipping capacity is only available at the annual frequency. We therefore construct a monthly series using the Chow-Lin temporal disaggregation method, implemented with the code suite of Quilis (2020). As monthly indicators, we use world industrial production and the oil price.

Figure 5: Impulse Responses to a Supply Chain Shock

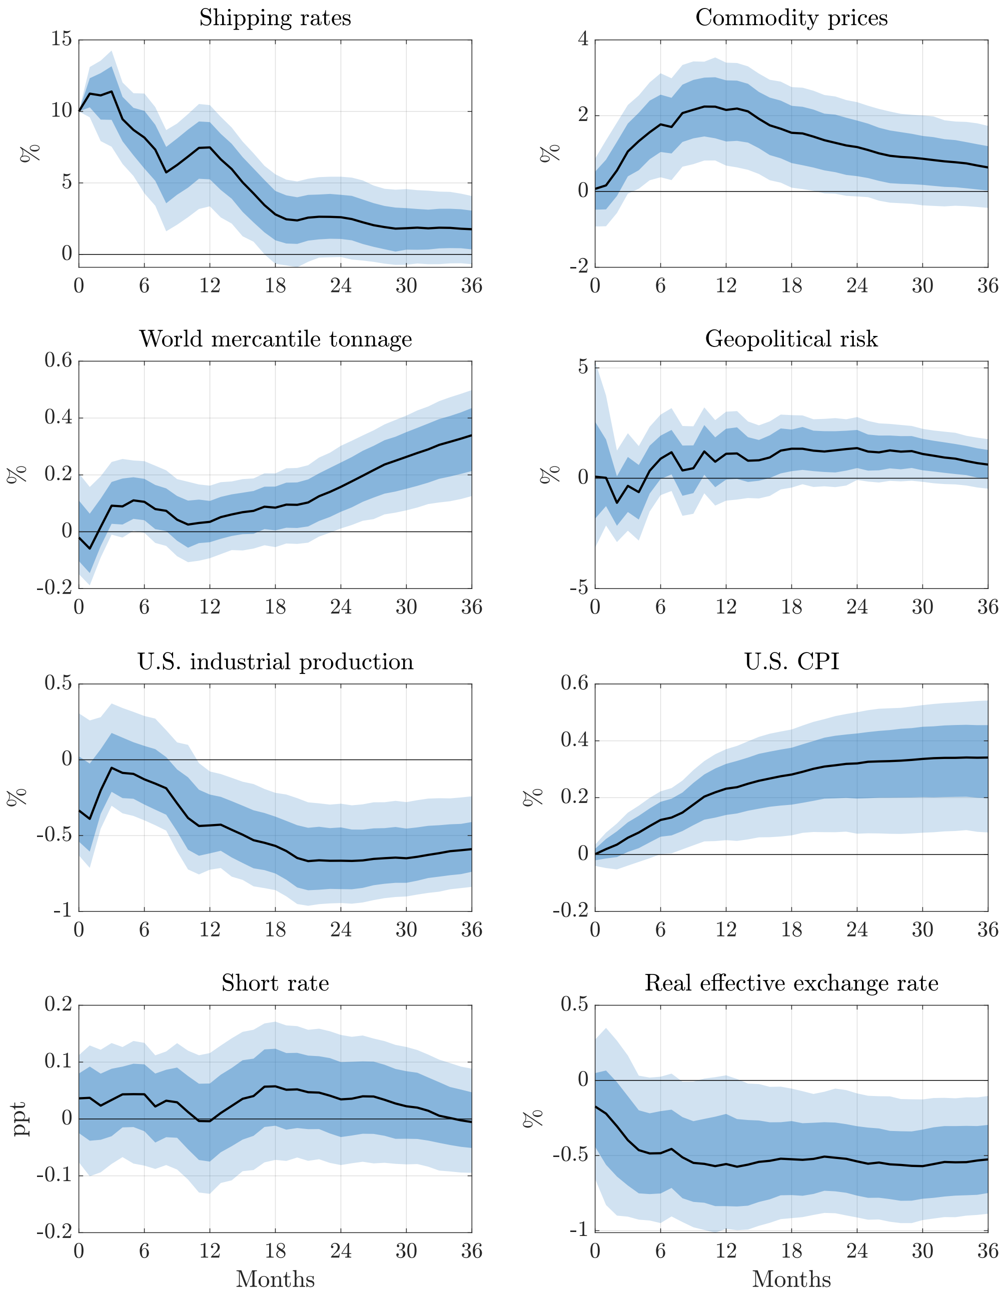

First stage: Robust F-statistic: 33.41, R^2: 3.46%, Adjusted R^2: 3.30%

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

reducing available capacity for scheduled shipments and creating bottlenecks at alternative ports, which disrupts the tightly coordinated global shipping network. These cascading effects ultimately impact global shipping capacity and pricing for months beyond the initial disruption.

The rise in shipping costs has important consequences for shipping markets and global commerce. Commodity prices increase significantly, although the peak effect materializes only with some lag. Shipping capacity, as measured by the world merchant fleet, initially contracts slightly, but the effect is small. Interestingly, capacity subsequently expands, likely in response to bottlenecks and persistently elevated shipping costs. Finally, the shock has no significant effect on geopolitical risk. This is reassuring, as we purposefully abstract from disruptive events related to geopolitical tensions, which may affect the global economy through channels unrelated to supply chains.

Higher shipping costs also put pressure on the economy. U.S. industrial production falls significantly and U.S. consumer prices increase persistently. The supply chain shock has stagflationary effects, creating a trade-off for monetary policy. This is reflected in the response of the short-term interest rate, which tends to increase even though the response is not statistically significant. Finally, the real effective exchange rate falls, implying a stark depreciation of the dollar.

Quantitatively, a supply chain shock that raises shipping rates by 10 percent increases commodity prices by about 2 percent and expands the world merchant fleet by 0.3 percent at its peak. U.S. industrial production falls by around 0.5 percent, the U.S. CPI increases by close to 0.4 percent, and the real effective exchange rate depreciates by 0.5 percent. These magnitudes are broadly comparable to those in Bini (2025), with two notable differences. The decline in industrial production is more persistent in our estimates, and commodity prices increase, whereas they fall in the short run in Bini (2025). This difference may reflect that our high-frequency design better accounts for demand-related confounding shocks.

Overall, the impacts are considerable but comparable to the economic effects of other supply-side shocks, such as commodity price shocks (e.g. Baumeister and Hamilton, 2019; Känzig, 2021). However, the sluggish increase in commodity prices, together with the delayed expansion of the world merchant fleet, suggests that we are not merely picking up commodity price shocks such as oil price shocks, which typically have more immediate effects on energy prices.

4.1. Addressing Identification Concerns

We now assess the robustness of our results to key assumptions underlying the empirical approach and discuss potential identification concerns related to the instrument.

First, we relax the invertibility requirement, estimating the responses based on invertibility-robust internal-instrument approach by Plagborg-Møller and Wolf (2019). We can see that the responses turn out to be very similar to our baseline external instruments results. Second, we assess the extent of lag truncation bias by estimating the dynamic responses based on local projections. The responses are very similar to our baseline responses, suggesting that our dynamic VAR structure is flexible enough to capture the dynamics in shipping markets and the U.S. economy adequately. Finally, we show results from a local projections–instrumental variables specification that relaxes both invertibility and the dynamic VAR structure, in Appendix C.2. Despite the power problem discussed above, the responses are similar to those from our baseline VAR, with industrial production adjusting somewhat more slowly.

We also address a number of additional concerns regarding the selection of our events. A first concern is that collisions at the maritime choke points could be more likely to occur during periods of economic expansion, potentially introducing some endogenous variation in our instrument. Reassuringly, our results turn out to be robust when excluding collision-related events, see Appendix Figure C.1.

Second, disruptions in the Suez Canal may coincide with major geopolitical developments, raising concerns about potential confounders. In our baseline specification, we already control for a news-based geopolitical risk index. To further mitigate these concerns, we show that our results are robust to only using events in the Panama Canal, see Appendix Figure C.7. Including the Suez Canal events yields similar point estimates but helps improve precision.

Third, we consider major macroeconomic episodes that could confound our estimates. We address the post-2019 disruptions by re-estimating the model on a pre-COVID sample, ending before the onset of the pandemic; this restriction also excludes the period around Russia's invasion of Ukraine. In addition, we control explicitly for the global financial crisis using a dummy variable. The responses remain close to the baseline, suggesting that these large macroeconomic episodes do not drive our results; see Appendix Figures C.12 and C.14.

Because our focus is on supply chain disruptions, positive shipping-cost surprises provide the most direct identifying variation. We therefore re-estimate the model using only positive surprises. The responses are very similar to the baseline, as shown in Appendix

Figure 6: Relaxing VAR Assumptions

(a) Invertibility-Robust Responses (b) Assessing Lag-Truncation Bias

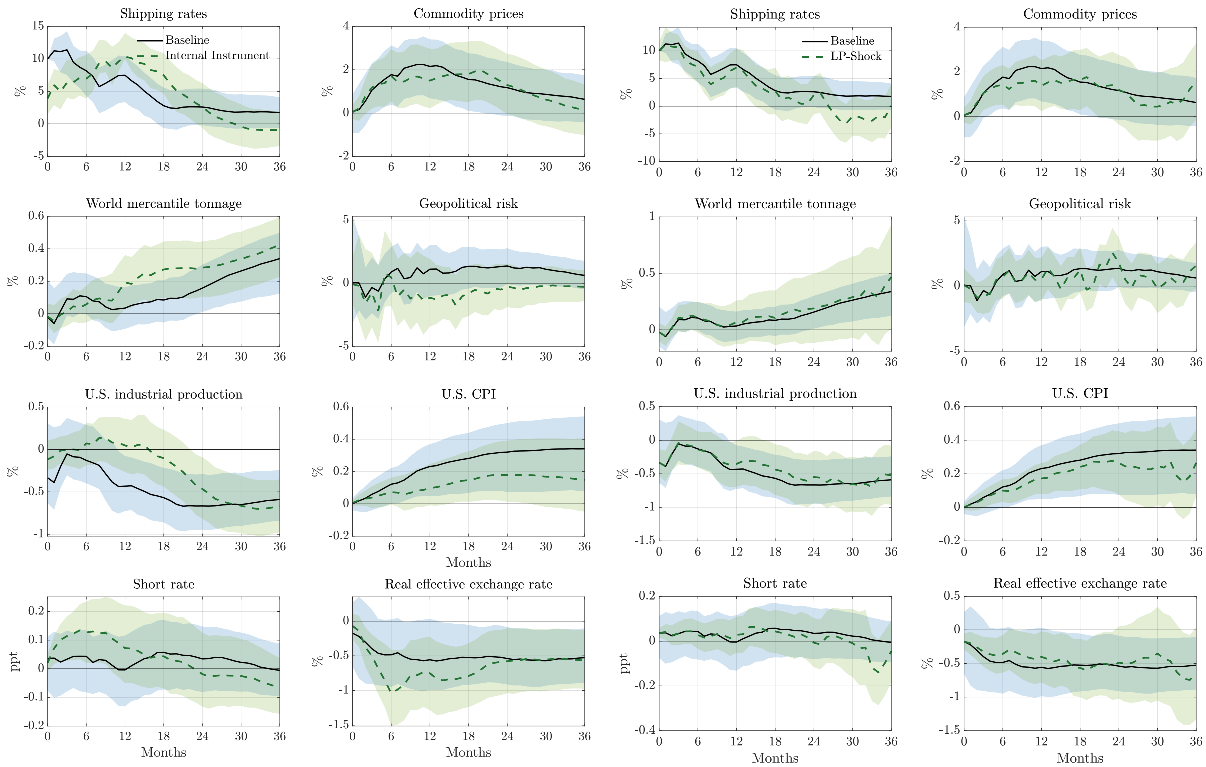

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. Panel (a) shows the responses under different invertibility-robust approaches. Panel (b) compares the responses estimated based on VARs and local projections. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

Figure C.8. We also show that the results are stable across alternative event windows (Appendix Figure C.6), which helps alleviate concerns that, for some events, our narrow baseline window may not capture the full market response.

We also explore how the predictability in our shipping cost surprises affects the results. We find consistent effects when restricting the sample to the first event in each sequence—those that are arguably least predictable—as well as when we forgo any attempt to remove predictability altogether, see Appendix Figures C.4 and C.5. This suggests that our findings are not mechanically driven by anticipatory behavior.

Finally, we show in Appendix C.3 that our results are robust to varying the sample period and other specification choices.

4.2. Wider Effects and Propagation Channels

To strengthen our interpretation of a supply chain shock, we study the effects of the shock on a wider range of financial macroeconomic variables. We follow the approach outlined in Section 3.1 to compute the impulse responses for additional variables of interest.

Impact on shipping costs. We begin by documenting that the identified shock leads to a broad-based increase in shipping costs across different segments of maritime transport. This is important for interpretation. Our baseline measure of shipping costs relies on the BDI, which captures dry bulk shipping rates. However, the disruptions we study—such as blockages at key maritime choke points—are not specific to bulk shipping and should affect the entire shipping network, including tanker and containerized freight markets.

Figure 7: Broader Effects on Shipping Costs

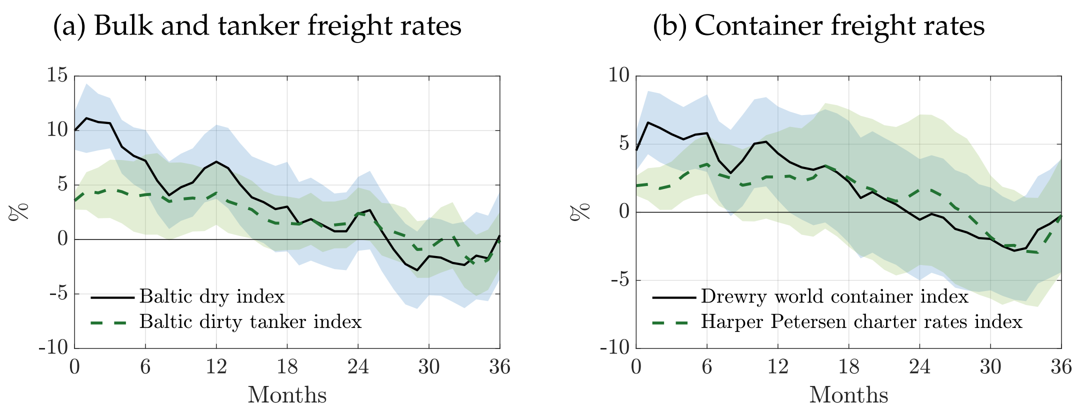

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The IRFs are estimated by local projections (10) on the aggregated supply chain shock extracted from our baseline external instruments VAR. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

To assess the generality of the shock, we extend our analysis to a range of alternative shipping cost indices covering tanker, charter, and containerized freight markets. Figure 7 shows that the supply chain shock induces a significant and persistent increase in shipping rates across all segments. In addition to the BDI, both clean and dirty tanker rates, charter rates, and container freight indices rise on impact and remain elevated for several months.

These results indicate that the identified shock captures a disruption to global shipping conditions rather than movements specific to dry bulk markets. Having established that the shock leads to a broad-based increase in shipping costs, we next turn to downstream manifestations of supply chain stress.

Shortages. To corroborate our interpretation of a supply chain shock, we first assess whether the shock is associated with significant delays in deliveries and shortages. To that end, we rely on the supplier delivery times indicator from the ISM and the newspaper-based supply chain shortage indices constructed by Caldara, Iacoviello, and Yu (2024). Figure 8 shows that the supply chain shock leads to a significant increase in delivery times and supply chain shortages—consistent with the notion that supply chain disruptions delay relevant shipments and cause shortages down the line.

Figure 8: Effects on Supply Chain Shortages and Delivery Times

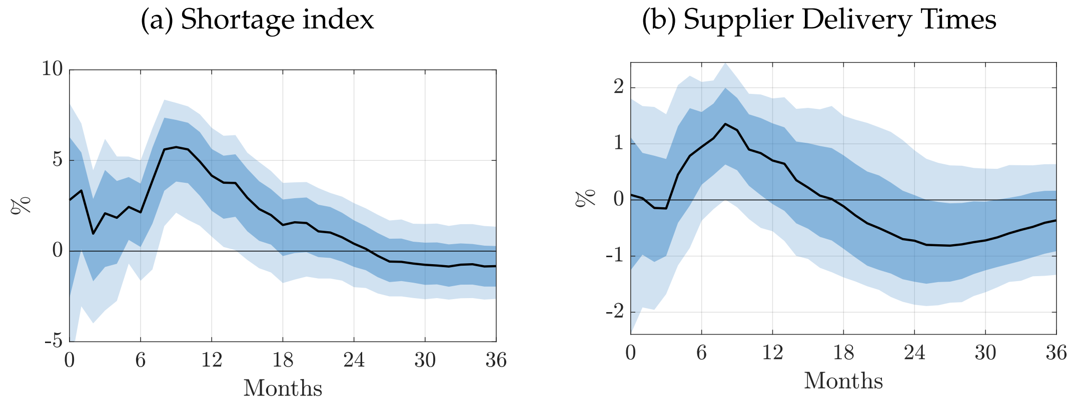

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The IRFs are obtained by using our baseline VAR and augmenting it by the variable of interest. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

Geopolitical risk. As shown above, the supply chain shock does not lead to a significant change in geopolitical risk. This suggests that our event selection, which excludes broad geopolitical events and geopolitically sensitive choke points such as the Strait of Hormuz, isolates variation in shipping conditions that is not driven by geopolitical risk. We also find no significant increase in other measures of risk and uncertainty, including financial uncertainty, economic and trade policy uncertainty, and crude oil volatility (Appendix Figure D.1).

Consumer prices and inflation expectations. We have seen that supply chain shocks lead to an increase in commodity and consumer prices. We now examine the components of consumer prices to identify which categories are most affected. In Figure 9, we show the responses of major consumer price components, including core, nondurable, energy, durable, and service prices, alongside the headline measure from our baseline specification. The responses of nondurable and energy prices qualitatively resemble the response of producer prices and exhibit a more immediate response compared to the headline measure. The price of services and durables also increase significantly, with the durable price response displaying substantial persistence. Consequently, we also document a sluggish but significant rise in core consumer prices.

Figure 9: Impact on Consumer Prices

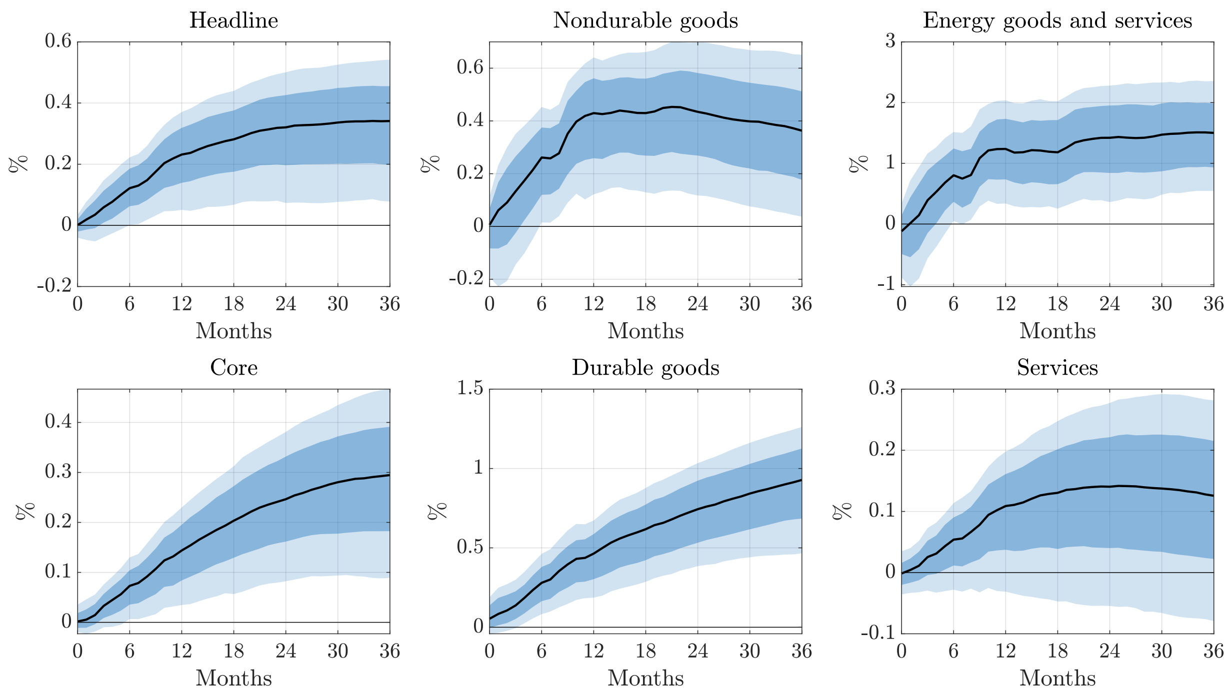

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The IRFs are obtained by using our baseline VAR and augmenting it by the variable of interest. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

Supply chain shocks may not only increase actual prices, but also translate into higher inflation expectations. Figure 10 shows the inflation expectations for consumers from the Michigan survey and for forecasters from the Survey of Professional Forecasters (SPF). Household inflation expectations increase significantly and persistently. Interestingly, the inflation expectations of professional forecasters do not respond.

Economic activity and trade. Supply chain disruptions have substantial effects on real activity and the external sector. Figure 11a shows that a supply chain shock leads to a broad-based decline in economic activity. Real GDP falls by about 0.4 percent at peak, with investment and consumption declining by around 1.5 percent and 0.3 percent, respectively. These responses are consistent with disruptions to production and higher in-

Figure 10: Inflation Expectations

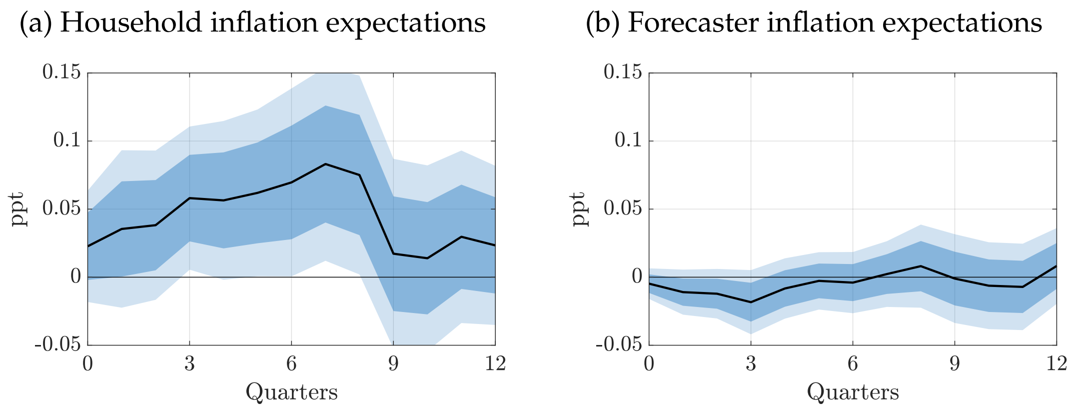

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The IRFs are obtained by using our baseline VAR and augmenting it by the variable of interest. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

put costs weighing on aggregate demand, especially investment.

Turning to external adjustment, the shock also has pronounced effects on trade. Figure 11b shows that the terms of trade deteriorate significantly, driven by a sharp increase in import prices. At the same time, the trade balance improves, consistent with the depreciation of the dollar and a contraction in imports. Together, these results highlight that supply chain shocks propagate through both domestic demand and international relative prices.

4.3. Understanding the Recent Inflationary Episode

How important were supply chain shocks in the 2021–2022 inflation surge? We use our identified supply chain shock to revisit this period and assess the quantitative contribution to macroeconomic dynamics. To this end, we simulate the economy under the sequence of estimated supply chain shocks while setting all other structural shocks to zero. This exercise provides a counterfactual estimate of how inflation, commodity prices, and real activity would have evolved in the absence of other shocks. The results are shown in Figure 12.

We find that supply chain disruptions contribute meaningfully to the variation in inflation over the 2020–2022 period. In particular, they help explain a substantial share of the rise in inflation through 2021, consistent with the timing of widespread logistical bottlenecks and maritime stress. However, they appear to contribute less later on. Our results suggest that while supply chain shocks were an important factor early on, they

Figure 11: Economic Activity and Trade Effects

(a) Macroeconomic Aggregates

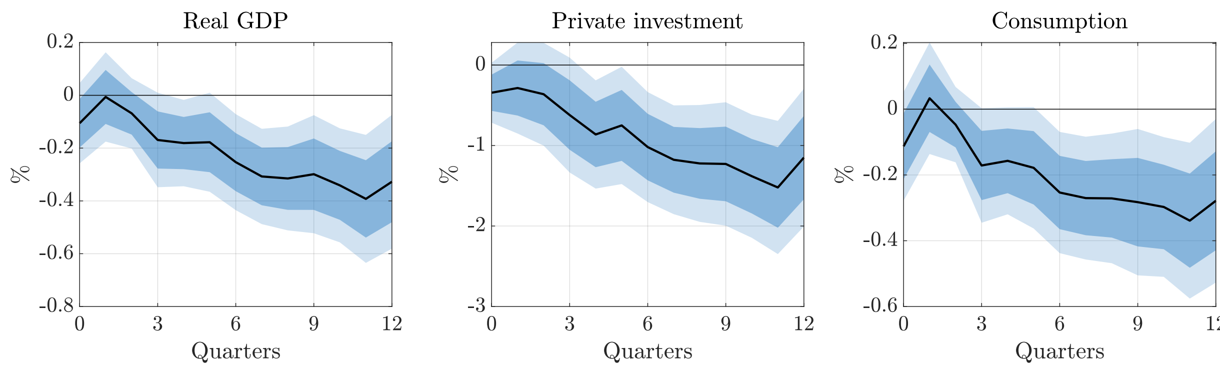

(b) Terms of Trade and Trade Balance

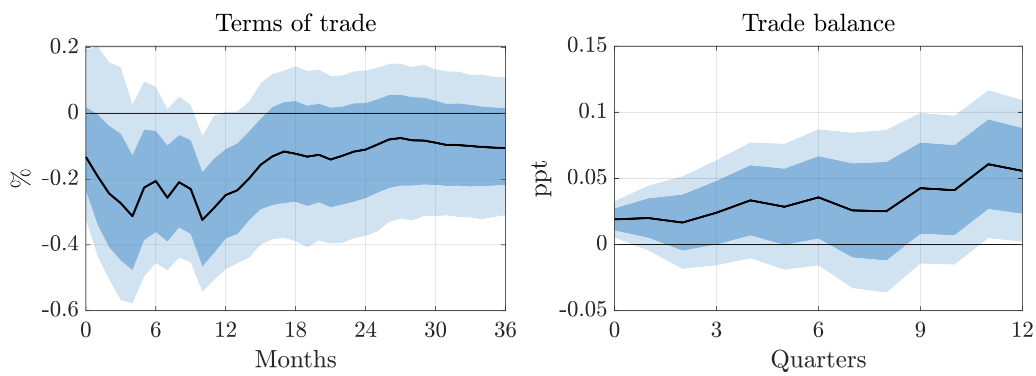

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The IRFs in Panel a are estimated by local projections (10) on the aggregated supply chain shock extracted from our baseline external instruments VAR. The IRFs in Panel b are obtained from our baseline VAR by augmenting the system with the variable of interest. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

cannot explain the sharp acceleration in inflation observed in 2022.

The same pattern is evident in commodity price dynamics. Supply chain disruptions explain a sizable portion of the fluctuations in commodity prices throughout 2021, reflecting their role in affecting global transportation costs and intermediate input availability. Yet, they fall short of accounting for the drastic commodity price spikes that followed the Russian invasion of Ukraine, pointing instead to additional geopolitical or supply-side factors unrelated to logistics.

Overall, we conclude that supply chain shocks played some role in the recent inflationary episode but cannot account for the peak inflation rates observed in 2022. This finding is in line with recent evidence highlighting the importance of expansionary demand-side policies—such as fiscal stimulus and accommodative monetary policy (Giannone and

Figure 12: The Role of Supply Chain Shocks in the 2021–2022 Inflation Surge

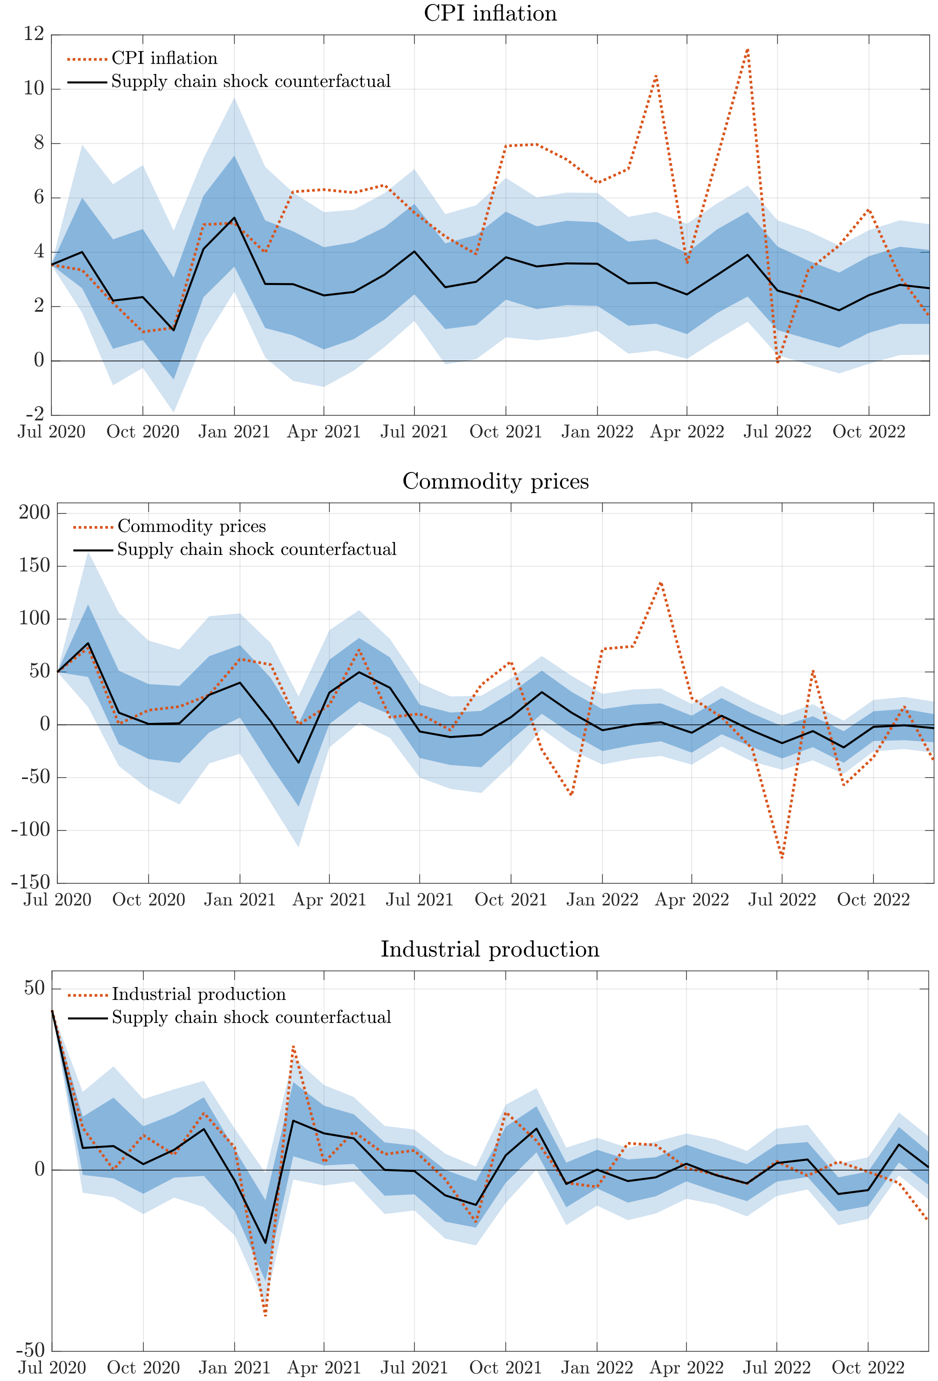

Notes: Counterfactual evolution of CPI inflation, commodity price inflation and industrial production growth when retaining only supply chain shocks from mid-2020 to 2022 (black line), together with the actual evolution of these variables (dashed red line). The dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

Primiceri, 2024)—as well as commodity supply shocks, particularly in energy and food markets (Gagliardone and Gertler, 2023).

That said, supply chain shocks had a more pronounced effect on real economic activity. Our analysis shows that they were an important drag on industrial production throughout 2022. This suggests that while supply-side bottlenecks had limited effects on the 2022 inflation peak, they significantly impaired production and contributed to the uneven sectoral recovery in the aftermath of the pandemic.

4.4. The Role of Monetary Policy

We now turn to the question of how monetary policy shapes the propagation of supply chain disruptions. To this end, we conduct a counterfactual policy simulation using the McKay and Wolf (2023) approach, which allows us to evaluate how supply chain shocks propagate under alternative monetary policy responses. The idea is to leverage estimated impulse responses to monetary policy shocks to impose a counterfactual monetary response to supply chain shocks. By only using a combination of shocks on impact, the contemplated counterfactual policy is incorporated in private-sector expectations ex-ante and thus robust to the Lucas critique. To implement the approach, we rely on the identified monetary policy shocks by Bauer and Swanson (2023).

This exercise allows us to quantify the potential trade-offs associated with a more aggressive monetary tightening aimed at offsetting the inflationary pressures induced by supply chain disruptions. Specifically, it sheds light on how costly it would be for the central bank to prevent the inflationary impacts of supply chain shocks by responding preemptively and forcefully to signals of supply chain stress.

The results, shown in Figure 13, highlight a key trade-off. A more aggressive monetary response can indeed stabilize prices following a supply chain shock by curbing inflationary pressures through substantial tightening of interest rates. However, this comes at substantial cost: the counterfactual path implies a much sharper contraction in industrial production, falling by almost 1.1 percent at peak.

Overall, these findings illustrate the complicated trade-offs monetary policy is confronted with in the face of pervasive supply chain pressures.

5. Sectoral Impacts and Substitution

Thus far, we have focused on the aggregate macroeconomic effects of supply chain shocks. These aggregate responses, however, may mask important heterogeneity in how

Figure 13: Monetary Policy Counterfactuals

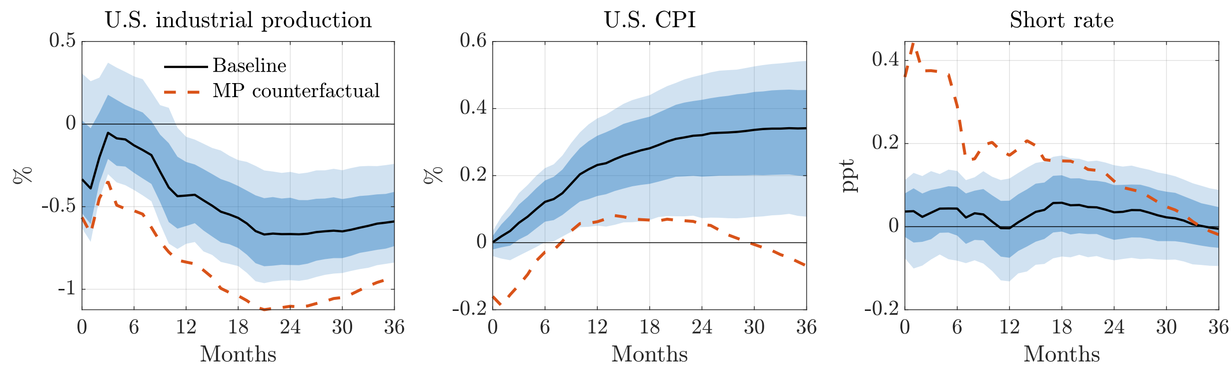

Notes: Responses of industrial production, CPI and the short-term interest rate from our baseline model (black line) against counterfactual responses when monetary policy aims to stabilize prices (dashed red line). The dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

shocks propagate across sectors. Studying this heterogeneity helps shed light on the transmission mechanisms at work and provides useful moments for disciplining quantitative production-network models. In this section, we document how sectoral output and prices respond to supply chain shocks and use these patterns to motivate our analysis of substitution across production inputs.

5.1. Sectoral Impacts

To better understand the transmission of supply chain shocks, we examine how their effects differ across sectors. We focus on sectoral responses in industrial production, distinguishing between non-durable and durable goods industries. This disaggregated perspective allows us to identify which parts of the production network are most affected by disruptions and how these effects vary across sectors with different input structures.

Figure 14 reports impulse responses for production and producer prices across a range of non-durable and durable goods production sectors. The effects of the supply chain shock are broad-based but display considerable heterogeneity across industries.

In non-durable goods sectors, production declines are particularly pronounced in industries that rely more heavily on intermediate inputs, such as chemicals and plastics, while other sectors such as food exhibit more muted responses. Price responses also vary notably across sectors. In particular, upstream industries such as petroleum and coal products display large and persistent increases in producer prices, whereas downstream sectors such as food and paper experience more moderate price pressures.

Durable goods sectors display even more pronounced heterogeneity in quantity re-

Figure 14: Sectoral Effects

Non-Durable Goods Production

(a) Industrial Production

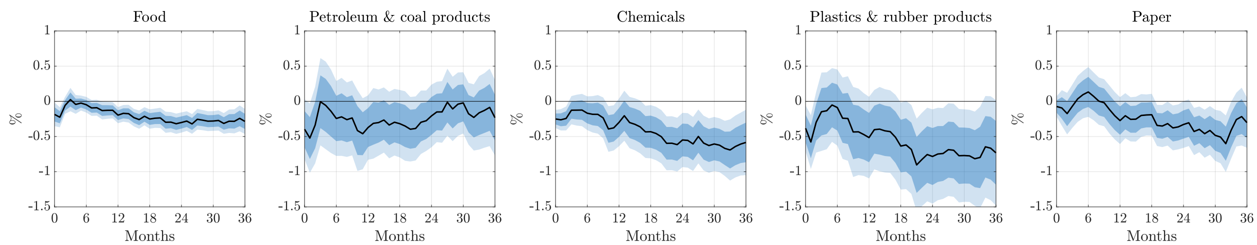

(b) Producer Prices

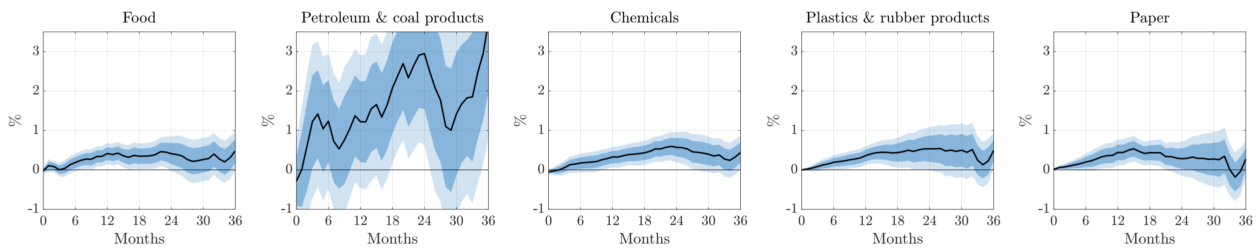

Durable Goods Production

(c) Industrial Production

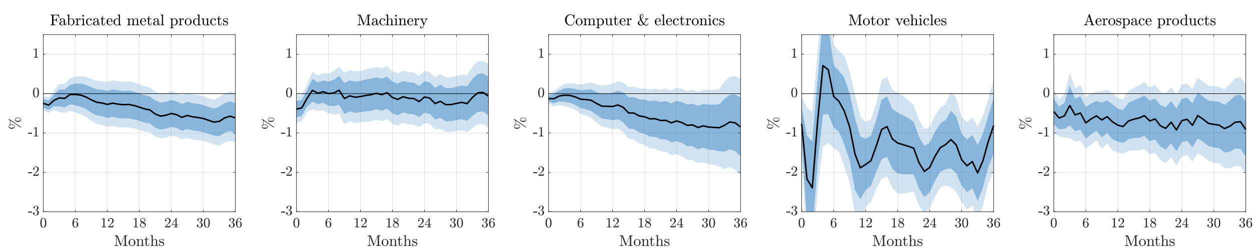

(d) Producer Prices

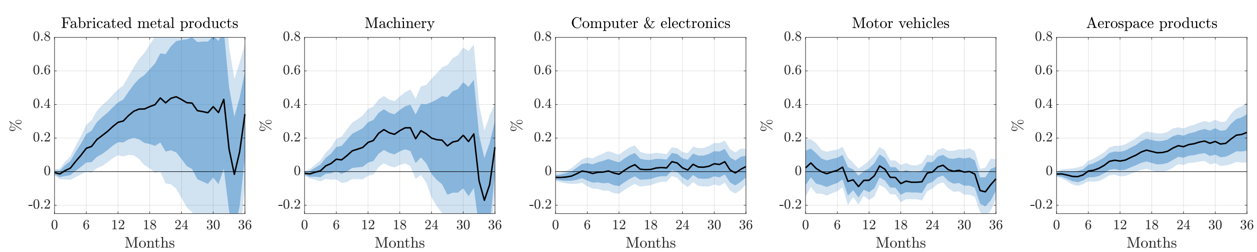

Notes: Impulse responses of non-durable and durable goods production sectors to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The IRFs are estimated by local projections (10) on the aggregated supply chain shock extracted from our baseline external-instruments VAR. The lines are point estimates, and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

sponses. The decline in production is concentrated in specific industries, most notably motor vehicles, with aerospace products also showing a sustained contraction. By contrast, sectors such as machinery and computer and electronic products exhibit more muted responses, although these estimates are imprecise and not statistically significant. Price responses in durable sectors are generally more moderate, with some increases in upstream segments such as metals, but less pronounced than in non-durable upstream industries.

Taken together, these patterns point to a heterogeneous transmission of supply chain shocks across sectors. Upstream industries experience strong price pressures, while quantity adjustments are concentrated in specific downstream sectors—particularly motor vehicles—and in input-intensive non-durable industries such as chemicals and plastics. These findings highlight the role of production networks in shaping the propagation of supply chain disruptions.

5.2. Estimating Substitution Elasticities

The sectoral evidence above points to substantial heterogeneity in how supply chain shocks propagate across industries. This heterogeneity is suggestive of differences in production structures and, in particular, in the scope for substituting across inputs when relative prices change. Motivated by these patterns, we now estimate elasticities of substitution across intermediate inputs. These estimates should not be interpreted as primitive structural production-function elasticities: the identified shock may affect input shares through several equilibrium channels beyond the relative price of a given input, including changes in other input prices, output demand, and income. We therefore interpret them as reduced-form macro elasticities that capture equilibrium reallocation across suppliers and sectors in response to supply chain shocks.

Framework. To guide the empirical specification, we consider a standard multi-sector production network model with a nested constant elasticity of substitution (CES) architecture, following Atalay (2017). Assume the economy consists of N perfectly competitive sectors. Firms in each sector combine capital, labor, and a composite bundle of intermediate inputs to produce gross output. In particular, output in sector i is given by

$$Y_{it} = A_{it} \left[(1-\mu_i)^{\frac{1}{\zeta}} V_{it}^{\frac{\zeta-1}{\zeta}} + \mu_i^{\frac{1}{\zeta}} M_{it}^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}},$$

where V_{it} is a capital-labor aggregate and M_{it} is a composite of intermediate inputs. The parameter ζ governs substitution between value added and intermediates.

The intermediate input bundle is itself a CES aggregate of inputs sourced from upstream sectors:

$$M_{it} = \left[\sum_{j=1}^{N} \omega_{ij}^{\frac{1}{\eta}} Y_{ijt}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where Y_{ijt} denotes inputs from sector j used by sector i, ω_{ij} captures the baseline importance of supplier j, and η is the elasticity of substitution across intermediate suppliers.

We focus on substitution within the intermediate bundle. Firms choose Y_{ijt} to minimize the cost of assembling M_{it}, taking input prices P_{jt} as given. The first-order condition for this problem implies:

$$P_{jt} = P_{it}^{M} \left(\frac{\omega_{ij} M_{it}}{Y_{ijt}} \right)^{\frac{1}{\eta}},$$

where P_{it}^{M} is the price index of the intermediate bundle. Rearranging yields the conditional demand for input j:

$$Y_{ijt} = \omega_{ij} M_{it} \left(\frac{P_{jt}}{P_{it}^{M}} \right)^{-\eta}.$$

Multiplying by prices and normalizing by total intermediate expenditure gives the expenditure share:

$$s_{ijt}^{M} \equiv \frac{P_{jt} Y_{ijt}}{P_{it}^{M} M_{it}} = \omega_{ij} \left(\frac{P_{jt}}{P_{it}^{M}} \right)^{1-\eta}.$$

Taking logs and first differences eliminates the time-invariant weights ω_{ij} and yields:

$$\Delta \ln s_{ijt}^{M} = (1-\eta) \Delta \ln \left(\frac{P_{jt}}{P_{it}^{M}} \right), \tag{11}$$

which relates changes in sourcing shares to changes in relative input prices. This is our key equation of interest. However, we will not estimate this equation directly using all observed variation in relative prices and expenditure shares. Instead, we focus on the responses of these variables to our identified supply chain shock, as described in detail below.

Data and empirical approach. Our estimation strategy requires three ingredients. First, we need time-varying information on intermediate input sourcing across industries to construct expenditure shares s^M_{ijt}. Second, we require measures of supplier-level output prices and purchaser-specific intermediate input price indices to form the relevant relative price term. Third, we use our identified supply chain shock, aggregated to the annual frequency. We construct expenditure shares using BEA Make and Use tables and combine these with BEA industry-level gross output deflators to build both output and intermediate input price indices. Appendix A.3 provides full details on the data construction.

Our final panel data set includes industry-pair observations spanning 1971 to 2022. For each pair (i, j), where i denotes a purchasing industry and j a supplying industry, we record the expenditure share s^M_{ijt}, the supplier's gross output price P_{jt}, and the purchaser-specific intermediate input price index P^M_{it}.

To estimate the elasticities, we proceed in two steps. In the first step, we estimate pair-specific responses of relative prices and expenditure shares to the annualized supply chain shock. For each pair (i, j), we estimate the contemporaneous response of relative prices:

$$\Delta \ln \left(\frac{P_{jt}}{P^M_{it}} \right) = \alpha^P_{ij} + \beta^P_{ij} SCShock_t + \varepsilon^P_{ijt}, \tag{12}$$

where β^P_{ij} captures how the relative cost of sourcing from supplier j responds to the shock.

Next, we estimate the response of expenditure shares at horizon h:

$$\ln s^M_{ij,t+h} - \ln s^M_{ij,t-1} = \alpha^Q_{ij}(h) + \beta^Q_{ij}(h) SCShock_t + \varepsilon^Q_{ijt}(h), \tag{13}$$

where $\beta^Q_{ij}(h)$ captures the dynamic adjustment in sourcing from supplier j by purchaser i.

In the second step, we relate the estimated sourcing-share responses to the corresponding relative-price responses across industry pairs, in cross-sectional regressions for each horizon h:

$$\widehat{\beta} ij^Q(h) = \mu_i(h) + \theta(h) \widehat{\beta} ij^P + u_{ij}(h), \tag{14}$$

where μ_i are purchasing-industry fixed effects. The inner-nest elasticity of substitution is then recovered as

$$\widehat{\eta}(h) = 1 - \widehat{\theta}(h). \tag{15}$$

By focusing on the responses of relative input prices and sourcing shares to identified supply chain disruptions, this procedure isolates variation in input prices driven by exogenous supply chain shocks and maps it into the equilibrium reallocation of sourcing across supplier-purchaser pairs. The resulting estimates are horizon-specific reduced-form macro elasticities: they distinguish between short-run and medium-run adjustment and capture reallocation across suppliers and sectors, rather than primitive production-function elasticities.

Elasticities of substitution. Figure 15 reports the estimated elasticity of substitution across intermediate inputs over the first three years following a supply chain shock. We show results both with and without purchasing-industry fixed effects, which absorb time-invariant differences in production structures across purchasers.

Figure 15: Elasticity of Substitution Across Intermediate Inputs

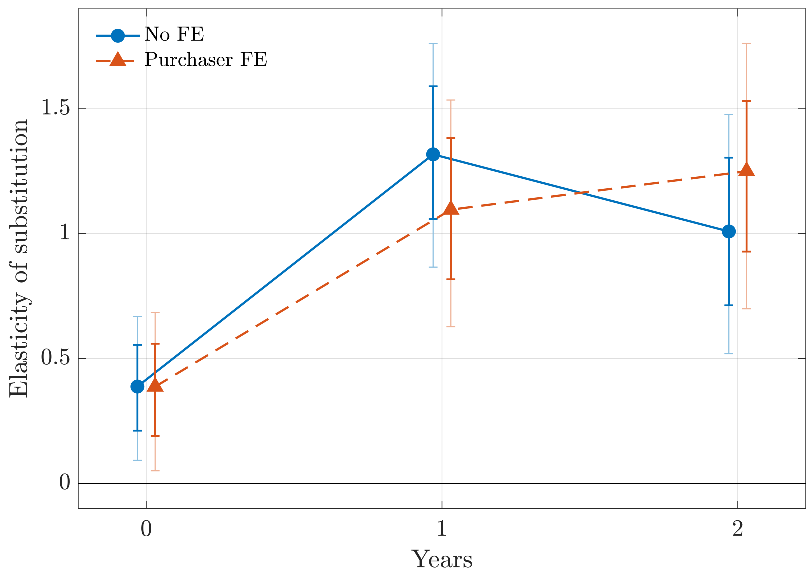

Notes: The figure reports horizon-specific estimates of the reduced-form macro elasticity of substitution between intermediate inputs, $\widehat{\eta}(h) = 1 - \widehat{\theta}(h)$. Estimates are obtained using the two-step local-projection procedure in equation (14), applied to the aggregated supply chain shock extracted from our baseline external-instruments VAR. Vertical bars denote 68 and 90 percent confidence intervals, respectively.

On impact, the estimated elasticity is small, around 0.4 across both specifications, in line with existing estimates in the literature (Atalay, 2017; Miranda-Pinto, 2021). This suggests that firms have little immediate scope to substitute away from a supplier that becomes relatively more expensive following a disruption. Such sluggish short-run adjustment is consistent with procurement rigidities, long-term supplier relationships, and the time required to identify and qualify alternative inputs.

After the first year, however, the elasticity rises sharply and remains elevated after two years. The medium-run estimates are slightly above one, although they are not statistically different from one. This indicates that substitution across suppliers takes place only gradually: firms appear to have limited scope to adjust sourcing patterns on impact, but reallocate more substantially over time as supply relationships adjust.

As a robustness check, Appendix D.2 estimates a one-step instrumental-variable specification directly in the industry-pair panel, instrumenting relative input prices with the supply chain shock interacted with predetermined supplier import exposure. In this specification, we can also include time fixed effects to better absorb confounding aggregate shocks. Reassuringly, the results are similar, suggesting that our baseline estimates are not driven by aggregate confounders.

Overall, the results imply that intermediate inputs are difficult to substitute in the short run but become meaningfully more substitutable at medium horizons. This dynamic pattern is important for quantitative production-network models, as it implies that the propagation of supply chain shocks depends not only on the static structure of input-output linkages, but also on the speed with which firms and sectors can reconfigure sourcing decisions.

6. Conclusion

In this paper, we provide new evidence on the macroeconomic implications of supply chain disruptions, exploiting the critical role of maritime trade choke points. Leveraging plausibly exogenous disruptive events at these bottlenecks in combination with high-frequency financial data, we identify structural supply chain shocks and trace their effects on the global economy. Our findings show that supply chain shocks have pervasive economic effects: they persistently raise shipping costs, commodity prices, and consumer prices, while depressing industrial production and broader economic activity. These stagflationary dynamics pose difficult challenges for monetary policy, as stabilizing prices in response to such shocks comes at the cost of a sharper decline in output.

Moving beyond aggregate outcomes, we document substantial heterogeneity in how these shocks propagate across sectors. Upstream industries experience strong price pressures, while output contractions are concentrated in downstream sectors such as motor vehicles and in input-intensive non-durable industries such as chemicals and plastics. Using granular input-output data, we estimate reduced-form macro elasticities of substitution across intermediate inputs. We find limited substitution in the short run, with elasticities well below one, but substantially stronger adjustment over time. After one year, the elasticity rises to slightly above one and remains elevated after two years, pointing to sluggish but economically meaningful reallocation of sourcing across suppliers and sectors.

Our results highlight the vulnerability of global supply chains and their substantial influence on macroeconomic outcomes. We show that supply chain shocks contributed meaningfully to the recent inflationary episode, especially through 2021, but cannot account for the sharp acceleration in inflation observed in 2022. As climate change, extreme weather, and geopolitical tensions increasingly disrupt global logistics, understanding the role of supply chain shocks is essential for designing appropriate policy responses.

References

Acemoglu, Daron, Ufuk Akcigit, and William Kerr (2016). "Networks and the Macroeconomy: An Empirical Exploration". *NBER Macroeconomics Annual* 30, pp. 273–335.

Acemoglu, Daron and Alireza Tahbaz-Salehi (2024). "The Macroeconomics of Supply Chain Disruptions". *Review of Economic Studies*, rdae038.

Alessandria, George, Shafaat Yar Khan, Armen Khederlarian, Carter Mix, and Kim J. Ruhl (2023). "The aggregate effects of global and local supply chain disruptions: 2020–2022". *Journal of International Economics* 146, p. 103788.

Atalay, Enghin (2017). "How important are sectoral shocks?" *American Economic Journal: Macroeconomics* 9.4, pp. 254–280.

Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, Ricardo Marto, and Francesco Zanetti (2025). *Shipping to America*. Working Paper. University of Oxford.

Bai, Xiwen, Jesús Fernández-Villaverde, Francesco Zanetti, and Yiliang Li (2024). "The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory". *Working Paper*.

Baqaee, David Rezza and Emmanuel Farhi (2019). "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem". *Econometrica* 87.4, pp. 1155–1203.

Barrot, Jean-Noël and Julien Sauvagnat (2016). "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks". *The Quarterly Journal of Economics* 131.3, pp. 1543–1592.

Bauer, Michael D. and Eric T. Swanson (2023). "A reassessment of monetary policy surprises and high-frequency identification". *NBER Macroeconomics Annual* 37.1, pp. 87–155.

Baumeister, Christiane and James D. Hamilton (2019). "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks". *American Economic Review* 109.5, pp. 1873–1910.

Benigno, Gianluca, Julian Di Giovanni, Jan J. Groen, and Adam I Noble (2022). "The GSCPI: A New Barometer of Global Supply Chain Pressures". *Federal Reserve Bank of New York Staff Reports* 1017.

Bini, Lapo (2025). "The Macroeconomic Effects of Global Supply Chain Shocks". *Available at SSRN 5375776*.

Blaum, Joaquin, Federico Esposito, and Sebastian Heise (2026). "Input Sourcing under Climate Risk: Evidence from U.S. Manufacturing Firms". *The Review of Economic Studies*, rdag031.

Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar (2019). "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake". *The Review of Economics and Statistics* 101.1, pp. 60–75.

Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar (2023). "The Long and Short (Run) of Trade Elasticities". *American Economic Review* 113.4, pp. 861–905.

Caldara, Dario and Matteo Iacoviello (2018). "Measuring geopolitical risk".

Caldara, Dario, Matteo Iacoviello, and David Yu (2024). "Measuring Shortages Since 1900". *Working Paper*.

Carrière-Swallow, Yan, Pragyan Deb, Davide Furceri, Daniel Jiménez, and Jonathan D. Ostry (2023). "Shipping costs and inflation". *Journal of International Money and Finance* 130, p. 102771.

Carvalho, Vasco M., Makoto Nirei, Yukiko U. Saito, and Alireza Tahbaz-Salehi (2021). "Supply chain disruptions: Evidence from the great east japan earthquake". *The Quarterly Journal of Economics* 136.2, pp. 1255–1321.

Carvalho, Vasco M. and Alireza Tahbaz-Salehi (2019). "Production Networks: A Primer". *Annual Review of Economics* 11.1, pp. 635–663.

Castro-Vincenzi, Juanma, Gaurav Khanna, Nicolas Morales, and Nitya Pandalai-Nayar (2024). *Weathering the storm: Supply chains and climate risk*. Tech. rep. National Bureau of Economic Research.

Cieslak, Anna (2018). "Short-rate expectations and unexpected returns in treasury bonds". *The Review of Financial Studies* 31.9, pp. 3265–3306.

Comin, Diego A., Robert C. Johnson, and Callum J. Jones (2023). *Supply chain constraints and inflation*. Tech. rep. National Bureau of Economic Research.

Di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A. Yildirim (2022). *Global supply chain pressures, international trade, and inflation*. Tech. rep. National Bureau of Economic Research.

Gagliardone, Luca and Mark Gertler (2023). "Oil prices, monetary policy and inflation surges".

Gertler, Mark and Peter Karadi (2015). "Monetary policy surprises, credit costs, and economic activity". *American Economic Journal: Macroeconomics* 7.1, pp. 44–76.

Giannone, Domenico and Giorgio Primiceri (2024). *The drivers of post-pandemic inflation*. Tech. rep. National Bureau of Economic Research.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson (2005). "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements". *International Journal of Central Banking* 1, pp. 55–93.

Herriford, Trenton, Elizabeth M Johnson, Nicholas Sly, A Lee Smith, et al. (2016). "How does a rise in international shipping costs affect US inflation?" *Macro Bulletin. Federal Reserve Bank of Kansas City*, pp. 1–3.

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

Jacks, David S. and Martin Stuermer (2021). "Dry bulk shipping and the evolution of maritime transport costs, 1850–2020". *Australian Economic History Review* 61.2, pp. 204–227.

Jentsch, Carsten and Kurt G. Lunsford (2019). "The dynamic effects of personal and corporate income tax changes in the United States: Comment". *American Economic Review* 109.7, pp. 2655–78.

Jordà, Òscar (2005). "Estimation and inference of impulse responses by local projections". *American Economic Review* 95.1, pp. 161–182.

Känzig, Diego R. (2021). "The macroeconomic effects of oil supply news: Evidence from OPEC announcements". *American Economic Review* 111.4, pp. 1092–1125.

— (2023). "The unequal economic consequences of carbon pricing". *Available at SSRN 3786030*.

Kuttner, Kenneth N. (2001). "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market". *Journal of Monetary Economics* 47.3, pp. 523–544.

Liu, Ernest and Aleh Tsyvinski (2020). *Dynamical structure and spectral properties of input-output networks*. Tech. rep. National Bureau of Economic Research.

McKay, Alisdair and Christian K. Wolf (2023). "What can time-series regressions tell us about policy counterfactuals?" *Econometrica* 91.5, pp. 1695–1725.

Mertens, Karel and Morten O. Ravn (2013). "The dynamic effects of personal and corporate income tax changes in the United States". *American Economic Review* 103.4, pp. 1212–47.

Minton, Robert and Brian Wheaton (2023). "Delayed Inflation in Supply Chains: Theory and Evidence". *Available at SSRN 4470302*.

Miranda-Agrippino, Silvia and Giovanni Ricco (2018). "Identification with external instruments in structural VARs under partial invertibility".

— (2021). "The Transmission of Monetary Policy Shocks". *American Economic Journal: Macroeconomics* 13.3, pp. 74–107.

Miranda-Pinto, Jorge (2021). "Production network structure, service share, and aggregate volatility". *Review of Economic Dynamics* 39, pp. 146–173.

Montiel Olea, José L., James H. Stock, and Mark W. Watson (2016). "Uniform inference in SVARs identified with external instruments".

Nakamura, Emi and Jón Steinsson (2018). "High-frequency identification of monetary non-neutrality: The information effect". *The Quarterly Journal of Economics* 133.3, pp. 1283–1330.

Peter, Alessandra and Cian Ruane (2025). "Intermediate Input Substitutability". *Working Paper*.

Plagborg-Møller, Mikkel and Christian K. Wolf (2019). "Local projections and VARs estimate the same impulse responses".

Quilis, Enrique M. (2020). *Temporal disaggregation*. https://www.mathworks.com/matlabcentral/fileexchange/69800-temporal-disaggregation.

Ramey, Valerie A (2016). "Macroeconomic shocks and their propagation". *Handbook of macroeconomics* 2, pp. 71–162.

Reuters (2004). *Suez Canal blocked by tanker, vessels back up*.

— (2015). *Panama temporarily restricts draft of vessels in the canal*.

Sims, Christopher A., James H. Stock, and Mark W. Watson (1990). "Inference in linear time series models with some unit roots". *Econometrica*, pp. 113–144.

Stock, James H. and Mark W. Watson (2012). "Disentangling the channels of the 2007-2009 recession". *Brookings Papers on Economic Activity*.

— (2018). "Identification and estimation of dynamic causal effects in macroeconomics using external instruments". *The Economic Journal* 128.610, pp. 917–948.

U.S. Energy Information Administration (2024). *Red Sea attacks increase shipping times and freight rates*. https://www.eia.gov/todayinenergy/detail.php?id=61363.

Online Appendix

Supply Chain Shocks and the Macroeconomy: Evidence from Global Shipping Disruptions

Diego R. Känzig[†]
Ramya Raghavan[‡]

Contents

[†]Northwestern University, CEPR and NBER. E-mail: dkaenzig@northwestern.edu. Web: www.diegokaenzig.com.

[‡]Northwestern University. E-mail: ramya.raghavan@u.northwestern.edu. Web: www.ramyaraghavan.com.

A. Data

A.1. Dataset on Events in Maritime Chokepoints

In this appendix, we provide additional information on the dataset of events that disrupted shipping traffic in critical maritime choke points.

We construct the event dataset from three sources. For both canals, we search newspaper archives through *Dow Jones' Factiva*, relying primarily on wire reports from *Reuters*, *Associated Press* (AP), and *Agence France-Presse* (AFP), which provide timely and consistently dated coverage of major disruptions such as large vessel incidents and weather-related disruptions. For the Panama Canal, we supplement these reports with the Panama Canal Authority's (ACP) "Advisory to Shipping" notices, available on the ACP's official website, which document operational restrictions such as draft limits. Finally, we use the Lloyd's List Intelligence serious-casualty records for the Suez and Panama Canal areas, which provide a structured vessel-level record of casualties, including groundings, collisions, machinery failures, contacts, and fires.

Our hand-collected Factiva–ACP inventory identifies 143 disruptions between 1970 and 2022, of which 45 occur in the Panama Canal and 98 in the Suez Canal. We document events that plausibly affected traffic through the relevant chokepoint and record the canal, type of disruption, description, severity, reporting date, and, when distinct, the date on which the disruption occurred. We then harmonize these events with Lloyd's serious-casualty records, which provide vessel-level information on maritime casualties, including vessel type, casualty date, primary casualty cause, and incident precis. After harmonization, Lloyd's contributes an additional 160 serious-casualty events to our dataset, of which 47 are in the Panama Canal area and 113 are in the Suez Canal area.[1]

A.2. Macro Data

In Table A.1, we provide information on the macroeconomic data used in the paper, including the data source and transformation used. The transformed series used in the baseline specification are depicted in Figure A.1.

[1]This procedure avoids double-counting disruptions that appear in both the hand-collected event inventory and the Lloyd's casualty records.

Table A.1: Data Description and Sources

Variable	Description	Source	Trans.
Instrument			
BDIY	Baltic dry index (daily average), extended using Hamilton (2021)	Bloomberg	$100 \times \Delta \log$
Baseline			
BDIY	Baltic dry index (daily average), extended using Hamilton (2021)	Bloomberg	$100 \times \log$
BCOM	Bloomberg commodity index (daily average)	Bloomberg	$100 \times \log$
TONNAGE	World merchant fleet (thousands of dead-weight tons; weight measure of a vessel's carrying capacity, including cargo, fuel, and stores), extended using Jacks and Stuermer (2021) and temporally disaggregated using world industrial production from Baumeister and Hamilton (2019) and oil prices from FRED (WTISPLC)	UNCTADStat	$100 \times \log$
GPRH	Historical geopolitical risk index from Caldara and Iacoviello (2022)	GPR index webpage	$100 \times \log$
INDPRO	Industrial production: total index	FRED	$100 \times \log$
PCEPI	Personal consumption expenditures: chain-type price index	FRED	$100 \times \log$
TB3MS	3-month treasury bill secondary market rate, discount basis	FRED	Level
RNUSBIS	Real narrow effective exchange rate for U.S.	FRED	$100 \times \log$
Shipping Rates			
BIDY	Baltic dirty tanker index (daily average), extended using BDIY	Bloomberg	$100 \times \log$
HARPEX	Harper Petersen shipping index (weekly average), extended using BDIY	Bloomberg	$100 \times \log$
WCIDCOMP	Drewry composite container freight benchmark rate (weekly average), extended using BDIY	Bloomberg	$100 \times \log$
Additional Monthly Variables			
SHORTAGE	Shortage index from Caldara, Iacoviello, and Yu (2024)	Shortage index webpage	$100 \times \log$
NAPMSUPL	Supplier delivery times, ISM manufacturing	Bloomberg	$100 \times \log$
VXO	CBOE S&P 100 volatility index, extended using S&P 500 from Bloomberg following Bloom (2009)	FRED	$100 \times \log$
EPU	Economic policy uncertainty index from Baker, Bloom, and Davis (2016)	Policy Uncertainty webpage	$100 \times \log$

Continued on next page

Table A.1: Data Description and Sources

Variable	Description	Source	Trans.
TPU	Trade policy uncertainty index from Caldara et al. (2020)	TPU index webpage	100 × log
OVXCLS	CBOE crude oil ETF volatility index extended using oil prices from FRED (WTISPLC) following Bloom (2009)	FRED	100 × log
USTOTPRCF	Terms of trade	Datastream	100 × log
Additional Quarterly Variables			
GDPC1	Real gross domestic product	FRED	100 × log
GPDIC1	Real gross private domestic investment	FRED	100 × log
PCECC96	Personal consumption expenditures	FRED	100 × log
TRADEBAL	Computed using real exports of goods and services (EXPGSC1), real imports of goods and services (IMPGSC1), and GDPC1 as: (EXPGSC1 - IMPGSC1)/GDPC1	FRED	Level
MICH	Median expected price change, 1-year horizon, Surveys of Consumers	FRED	Level
CPI6	Median expected change in CPI level, 1-year horizon, Survey of Professional Forecasters	Philadelphia Fed's webpage	Level
Consumer Prices			
DNDGRG-3M086SBEA	Personal consumption expenditures: nondurable goods (chain-type price index)	FRED	100 × log
DNRGRG-3M086SBEA	Personal consumption expenditures: energy goods and services (chain-type price index)	FRED	100 × log
PCEPILFE	Personal consumption expenditures excluding food and energy (chain-type price index)	FRED	100 × log
DDURRG-3M086SBEA	Personal consumption expenditures: durable goods (chain-type price index)	FRED	100 × log
DSERRG-3M086SBEA	Personal consumption expenditures: services (chain-type price index)	FRED	100 × log
Producer Prices			
PCU311311	Producer price index by industry: food manufacturing	FRED	100 × log
PCU324324	Producer price index by industry: petroleum and coal products manufacturing	FRED	100 × log
PCU325325	Producer price index by industry: chemical manufacturing	FRED	100 × log

Continued on next page

Table A.1: Data Description and Sources

Variable	Description	Source	Trans.
PCU326326	Producer price index by industry: plastics and rubber products manufacturing	FRED	$100 \times \log$
PCU322322	Producer price index by industry: paper manufacturing	FRED	$100 \times \log$
PCU332332	Producer price index by industry: fabricated metal product manufacturing	FRED	$100 \times \log$
PCU333333	Producer price index by industry: machinery manufacturing	FRED	$100 \times \log$
PCU334334	Producer price index by industry: computer and electronic product manufacturing	FRED	$100 \times \log$
PCU33613361	Producer price index by industry: motor vehicle manufacturing	FRED	$100 \times \log$
PCU33643364	Producer price index by industry: aerospace product and parts manufacturing	FRED	$100 \times \log$
Industrial Production			
IPFPNSS	Industrial production: final products and nonindustrial supplies	FRED	$100 \times \log$
IPMAT	Industrial production: materials	FRED	$100 \times \log$
IPG311S	Industrial production: manufacturing: nondurable goods: food	FRED	$100 \times \log$
IPG324S	Industrial production: manufacturing: nondurable goods: petroleum and coal products	FRED	$100 \times \log$
IPG325S	Industrial production: manufacturing: nondurable goods: chemical	FRED	$100 \times \log$
IPG326S	Industrial production: manufacturing: nondurable goods: plastics and rubber products	FRED	$100 \times \log$
IPG322S	Industrial production: manufacturing: nondurable goods: paper	FRED	$100 \times \log$
IPG332S	Industrial production: manufacturing: durable goods: fabricated metal product	FRED	$100 \times \log$
IPG333S	Industrial production: manufacturing: durable goods: machinery	FRED	$100 \times \log$
IPG334S	Industrial production: manufacturing: durable goods: computer and electronic product	FRED	$100 \times \log$
IPG3361S	Industrial production: manufacturing: durable goods: motor vehicle	FRED	$100 \times \log$
IPG3364S	Industrial production: manufacturing: durable goods: aerospace product and parts	FRED	$100 \times \log$

Figure A.1: Transformed Data Series

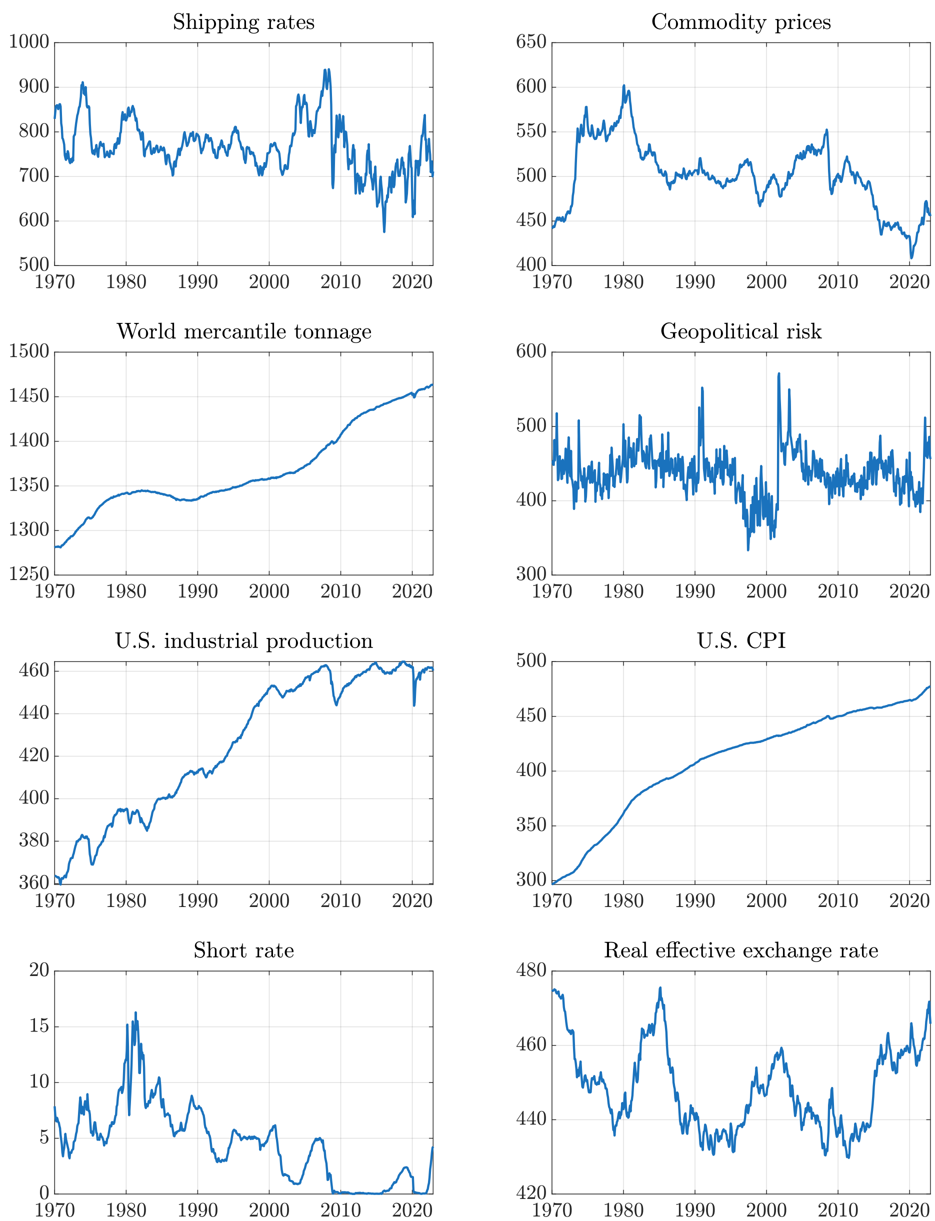

A.3. Sectoral Data

A.3.1. Input-Output Data

Make and Use tables. We rely on the BEA Make and Use tables at the annual frequency from 1970 to 2022. The *Make table* records, for each industry, the dollar value of each commodity it produces. Its rows are industries and its columns are commodities, so the (j, c) entry is the value of commodity c produced by industry j. Because industries often produce secondary commodities alongside their primary commodity output—for example, a steel mill could also sell scrap metal—the Make table is generally not diagonal.

The *Use table* records, for each industry, the dollar value of each commodity it purchases as an intermediate input, as well as its value added. Its rows are commodities and its columns are industries, so the (c, i) entry is the value of commodity c used by industry i. The Use table also includes final demand (consumption, investment, government, and net exports) as additional columns. Together, the Make and Use tables provide a complete account of how commodities flow between producers and users in the economy.

We also use the BEA *Import table*, which records, for each industry, the dollar value of imported commodities used as intermediate inputs. The structure mirrors the intermediate block of the Use table: rows are commodities and columns are industries, so the (c, i) entry is the value of imported commodity c used by industry i. We sum across commodity rows to obtain total imported intermediate inputs for each industry, which we use to compute the share of intermediate inputs sourced from abroad.

Redefinitions. The BEA publishes each of the above tables in two versions. The *before-redefinitions* tables record each establishment's entire output—primary and secondary commodities alike—under the industry defined by the establishment's primary activity. The *after-redefinitions* tables apply a secondary-product transfer: output of a commodity that is not the primary product of an industry is reassigned ("redefined") to the industry for which that commodity is primary. For example, a hotel that operates an on-site restaurant generates food and beverage revenue as a secondary product alongside its primary accommodation output. In the before-redefinitions table, revenue from both accommodation and food services is recorded under the accommodation industry. In the after-redefinitions table, the food and beverage portion is transferred to the food services industry, for which it is the primary product.

We use before-redefinitions tables for two reasons. First, consistency: redefinitions change over time, so a before-redefinitions panel maintains a stable commodity classification from 1970 to 2022. Second, alignment with our price data: we rely on BEA industry-level gross output price deflators, which are constructed on an establishment basis and therefore better align with the before-redefinition concept of industry output (Horowitz and Planting, 2009, p. 4–5).

Producer vs. purchaser prices. The Use tables are published in two valuations: at producer prices, which record what the seller receives, and at purchaser prices, which add wholesale and retail trade margins and transport costs to obtain what the buyer pays. We use the producer-price version throughout, as this is the only valuation available for the historical 1970 to 1996 tables.

Industry coverage. The 1970–1996 tables are defined over 65 industries; the 1997–2022 tables over 71. We harmonize the post-1997 tables to the 65-industry classification using the standard consolidations: retail trade subcategories `441`, `445`, `452`, and `4A0` to `44RT`; real estate subcategories `HS` and `ORE` to `531`; hospitals, `622`, and nursing and residential care facilities, `623`, to `622HO`; federal general government defense, `GFGD`, and non-defense, `GFGN`, to `GFG`. Next, following Atalay (2017), we further aggregate the BEA summary-level industries, as defined in Table A.2.

Notation. Let V denote the $N \times N$ Make matrix (industries $\times$ commodities), U the $N \times N$ intermediate portion (i.e., excluding final demand columns and value-added rows) of the Use matrix (commodities $\times$ industries), q the $N \times 1$ vector of total commodity output, and g the $N \times 1$ vector of total industry output.

A.3.2. Construction of Expenditure Shares

Our object of interest is the nominal expenditure share s^M_{ijt}: the fraction of industry i's total intermediate goods expenditure allocated to inputs sourced from industry j in year t. Because BEA tables record commodity flows rather than industry-to-industry flows, we follow Horowitz and Planting (2009) (p. 12-21) to redistribute commodity purchases back to their producing industries. The construction proceeds in three steps.

Step 1: Market shares. Define the market-share matrix as:

$$D = V\hat{q}^{-1},$$

where $\hat{q}$ is the diagonal matrix with total commodity output q on the diagonal. The (j, c) entry $D_{jc} = V_{jc}/q_c$ is the share of commodity c's output produced by industry j. Each

Table A.2: Aggregation of BEA Summary-Level Industries

Sector	BEA Summary Codes
Agriculture & Forestry	111CA, 113FF
Oil & Gas Extraction	211, 213
Mining	212
Utilities	22
Construction	23
Food & Kindred Products	311FT
Textile Mills	313TT
Apparel & Leather Products	315AL
Wood Products	321
Paper & Allied Products	322
Printing & Publishing	323, 511
Petroleum & Coal Products	324
Chemical Products	325
Plastics & Rubber Products	326
Non-Metallic Mineral Products	327
Primary Metals	331
Fabricated Metal Products	332
Machinery	333
Computer & Electronic Products	334
Electrical Equipment	335
Motor Vehicles	3361MV
Other Transport Equipment	3364OT
Furniture	337
Misc. Manufacturing	339
Wholesale & Retail Trade	42, 44RT
Transportation & Warehousing	481, 482, 483, 484, 485, 486, 487OS, 493
Communications	512, 513, 514
FIRE	521CI, 523, 524, 525, 531, 532RL
Other Services	5411, 5412OP, 5415, 55, 561, 562, 61, 621, 622HO, 624, 711AS, 713, 721, 722, 81
Government	GFG, GFE, GSLG, GSLE

Notes: Aggregation follows Atalay (2017), Table 7. Codes are BEA summary-level industry codes from the Input-Output Accounts. 44RT consolidates retail trade sub-categories (441, 445, 452, 4A0); 531 consolidates housing services and other real estate (HS, ORE); 622HO consolidates hospitals and nursing/residential care (622, 623); GFG consolidates federal general government defense and nondefense (GFGN, GFGD).

column of D sums to one.

Step 2: Industry-by-industry nominal use. Pre-multiplying the Use matrix by D converts commodity-based flows into industry-to-industry flows:

$$\widetilde{U} = D\,U. \tag{1}$$

The (j, i) entry $\sum_c D_{jc}\,U_{ci}$ is the total dollar value of purchases by industry i attributable to producing industry j.

Step 3: Expenditure shares. Finally, we normalize each column of $\widetilde{U}$ by total intermediate spending to obtain expenditure shares:

$$s^M_{ijt} = \frac{\tilde{U}_{ji}}{\sum_k \widetilde{U}_{ki}}. \tag{2}$$

Each column sums to one by construction.

Numerical illustration. We illustrate each step using the sample Make and Use tables from Horowitz and Planting (2009) (Chapter 12).

The Make table (V), with total industry output in the final column, is:

Table A.3: Make Table

	Commodity				
Industry	A	B	C	Scrap	Total industry output
A	300	25	0	3	328
B	30	360	20	2	412
C	0	15	250	0	265
Total commodity output	330	400	270	5	

The Use table, including value added and final demand, is:

Table A.4: Use Table

Commodity	Industry A	B	C	Final demand	Total commodity output
A	50	120	120	40	330
B	180	30	60	130	400
C	50	150	50	20	270
Scrap	1	3	1	0	5
Value added	47	109	34		190
Total industry output	328	412	265	190	

The intermediate block U consists of the first three commodity rows and the three industry columns (excluding the scrap row, value added, and final demand).[2]

Step 1. Dividing each column of the Make table V (excluding the scrap column) by total commodity output q gives the market-share matrix D:

Table A.5: Market Shares Matrix

Industry	Commodity A	B	C
A	$300/330 = 0.909$	$25/400 = 0.063$	$0/270 = 0.000$
B	$30/330 = 0.091$	$360/400 = 0.900$	$20/270 = 0.074$
C	$0/330 = 0.000$	$15/400 = 0.038$	$250/270 = 0.926$
Total commodity output	1.000	1.000	1.000

For instance, industry A produces 90.9% of all commodity A output and 6.3% of commodity B output.

[2] In practice, when computing our baseline expenditure shares, we include scrap as an intermediate input commodity. This is to ensure consistency with the latest available BEA methodology for industry-by-industry total requirements tables, which are constructed from the complete Use matrix, including scrap rows. In practice, scrap purchases are small relative to total intermediate spending for most industries, and our results are insensitive to this choice.

Step 2. Pre-multiplying U by D converts commodity-by-industry flows into industry-by-industry flows. For example:

$$\widetilde{U}_{AA} = (0.909 \times 50) + (0.0625 \times 180) + (0.000 \times 50) = 56.70.$$

Computing all nine entries gives the full industry-by-industry nominal use matrix $\widetilde{U}$:

Table A.6: Industry-by-Industry Use Table

	Purchasing Industry		
Supplying Industry	A	B	C
A	56.70	110.97	112.84
B	170.25	49.02	68.61
C	53.05	140.01	48.55
Total intermediate spending	280.00	300.00	230.00

Step 3. Dividing each column by its total yields the expenditure shares:

Table A.7: Industry-by-Industry Intermediate Expenditure Shares

	Purchasing Industry		
Supplying Industry	A	B	C
A	$56.70/280 = 0.203$	$110.97/300 = 0.370$	$112.84/230 = 0.491$
B	$170.25/280 = 0.608$	$49.02/300 = 0.163$	$68.61/230 = 0.298$
C	$53.05/280 = 0.189$	$140.01/300 = 0.467$	$48.55/230 = 0.211$
Column sum	1.000	1.000	1.000

Each entry (j, i) is the share of industry i's total intermediate budget allocated to supplier j. For instance, industry A allocates 60.8% of its intermediate budget to inputs from industry B and 20.3% to inputs from industry A itself. These shares form the time-varying weights entering our estimating equation (11).

A.3.3. Price Data

Source. We construct annual output and intermediate input price indices drawing on the BEA *GDP by Industry* accounts. We use the chain-type price indices for gross output for the period 1970 to 2022.[3] We also rely on nominal gross output values for constructing aggregation weights over the same period.

Splicing modern and historical series. The modern (1997 to 2022) and historical (1970 to 1997) price indices share a common year at 1997. We thus splice the two series by rebasing the historical index onto the modern scale at 1997. Thus, for each industry j, we multiply each historical observation by the ratio of the modern to historical index value in 1997,

$$\widetilde{P}_{jt} = P^{\text{hist}}_{jt} \times \frac{P^{\text{modern}}_{j,1997}}{P^{\text{hist}}_{j,1997}}, \qquad t < 1997,$$

so that the spliced series is expressed in the same units as the modern data throughout (with a base year of 2017).

A.3.4. Construction of Output and Intermediate Input Price Indices

Gross output price indices. The price indices for the 65 BEA summary-level industries are aggregated to the broader sectors using a Törnqvist price index following Atalay (2017). For each sector i and adjacent years $t-1$ and t, the log change in the price index is:

$$\Delta \log P_{i,t} = \sum_{j \in S} \overline{w}_{j,t} \, \Delta \log P_{j,t},$$

where S denotes summary-level industries part of the broader sector i and the Törnqvist weight is the average nominal gross output share across the two years,

$$\overline{w}_{j,t} = \frac{1}{2}\left(s_{j,t-1} + s_{j,t}\right), \qquad s_{j,t} = \frac{\text{Output}_{j,t}}{\sum_{k \in S} \text{Output}_{k,t}}.$$

Log changes are chained from 1970 to 2022, and the resulting log-level index is rebased so that each sector's price equals 100 in 2017.

[3]The combined hospital and nursing care industry (622HO) is reported as two separate lines in the modern files—Hospitals (line 78) and Nursing and Residential Care Facilities (line 79)—but as a single combined line 78 in the historical files. For the modern period, we construct a single 622HO price index by using the Törnqvist aggregation. For the historical period, line 78 is used directly.

Intermediate input price indices. For each sector i and adjacent years $t-1$ and t, the log change in the intermediate input price index is:

$$\Delta \log P_{i,t}^{M} = \sum_{j=1}^{30} \overline{w}_{ij,t}^{M} \, \Delta \log P_{j,t},$$

where $P_{j,t}$ is the gross output price index for supplier sector j and the Törnqvist weight is the average intermediate input expenditure share across adjacent years,

$$\overline{w}_{ij,t}^{M} = \frac{1}{2}\left(s_{ij,t-1}^{M} + s_{ij,t}^{M}\right).$$

The expenditure shares $s_{ij,t}^{M}$ are constructed as described in Equation 2. The log changes are chained and the resulting index is rebased to 2017.

B. Diagnostics of the Surprise Series

As discussed in the paper, we perform a number of additional checks to assess the validity of the surprise series. In this section, we investigate the autocorrelation and forecastability of the surprise series, as well as its correlation with other shocks from the literature.

Autocorrelation. Figure B.1 reports the autocorrelation function of the surprise series. The series exhibits some serial correlation, likely reflecting the clustering of certain event types. As a robustness check, we show that our results are virtually unchanged after orthogonalizing the surprise series with respect to its own lags, following Miranda-Agrippino and Ricco (2021) (see Figure C.5).

Figure B.1: Autocorrelation Function

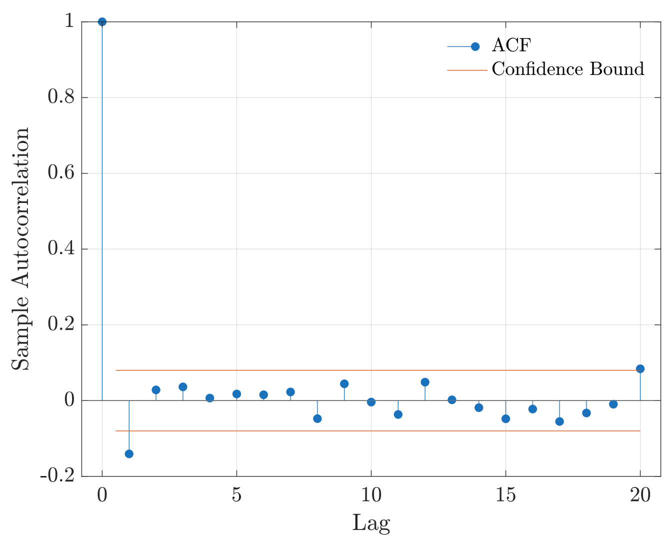

Granger causality tests. In Table B.1, we present the results of a number of Granger causality tests. We find little evidence that macroeconomic or financial variables have any power in forecasting the surprise series. For all variables considered, the p-values for the Granger causality test are far above conventional significance levels, with the joint test having a p-value of 0.19.

Table B.1: Granger Causality Tests

Variable	p-value
Instrument	0.1123
Shipping rates	0.3744
Commodity prices	0.1632
World mercantile tonnage	0.7117
Geopolitical risk	0.3967
U.S. industrial production	0.2874
U.S. CPI	0.3527
Short rate	0.1526
Real effective exchange rate	0.5532
Oil price	0.8579
VIX	0.9655
Shortage index	0.3557
Joint	0.1877

Notes: This table shows the p-values of a series of Granger causality tests of the shipping cost surprise series using a selection of macroeconomic and financial variables. To be able to conduct standard inference, the series are made stationary by taking first differences where necessary. The lag order is set to 6 and in terms of deterministics, only a constant term is included.

Correlation with other shock measures. Finally, in Table B.2 we examine whether the shipping costs surprise series is correlated with other shocks from the literature. The surprise series is uncorrelated with other structural shock measures from the literature, including oil, productivity, news, monetary policy, uncertainty, financial, and fiscal policy shocks.

Table B.2: Correlation with Other Shock Measures

Shock	Source	ρ	p-value	n	Sample
Panel A: Oil shocks					
Oil supply	Baumeister and Hamilton (2019)	0.04	0.44	456	1985M01-2022M12
	Caldara, Cavallo, and Iacoviello (2019)	-0.02	0.69	372	1985M01-2015M12
	Känzig (2021)	-0.02	0.71	456	1985M01-2022M12
Panel B: Productivity shocks					
Productivity	Basu, Fernald, and Kimball (2006)	0.01	0.96	108	1985Q1-2011Q4
	Smets and Wouters (2007)	-0.08	0.51	80	1985Q1-2004Q4
Panel C: News shocks					
News	Barsky and Sims (2011)	-0.06	0.55	91	1985Q1-2007Q3
	Beaudry and Portier (2014)	0.11	0.26	111	1985Q1-2012Q3
	Kurmann and Otrok (2013)	0.06	0.59	82	1985Q1-2005Q2
Panel D: Monetary policy shocks					
Monetary policy	Bauer and Swanson (2023)	0.00	0.92	383	1988M02-2019M12
	Gertler and Karadi (2015)	0.02	0.73	324	1990M01-2016M12
	Romer and Romer (2004)	-0.02	0.80	144	1985M01-1996M12
	Smets and Wouters (2007)	0.09	0.45	80	1985Q1-2004Q4
Panel E: Uncertainty and risk shocks					
VIX	Bloom (2009)	-0.01	0.84	456	1985M01-2022M12
Economic policy uncertainty	Baker, Bloom, and Davis (2016)	-0.00	0.93	450	1985M07-2022M12
Geopolitical risk	Caldara and Iacoviello (2022)	0.01	0.79	456	1985M01-2022M12
Panel F: Financial shocks					
Financial	Bassett et al. (2014)	-0.07	0.53	76	1992Q1-2010Q4
	Gilchrist and Zakrajšek (2012)	0.03	0.53	456	1985M01-2022M12
Panel G: Fiscal policy shocks					
Fiscal policy	Fisher and Peters (2010)	-0.13	0.19	96	1985Q1-2008Q4
	Ramey (2011)	-0.03	0.76	104	1985Q1-2010Q4
	Romer and Romer (2010)	-0.11	0.31	92	1985Q1-2007Q4

Notes: This table shows the correlation of the shipping cost surprise series with a wide range of structural shock measures from the literature. ρ is the Pearson correlation coefficient, the p-value corresponds to the test whether the correlation is different from zero, and n is the sample size. When the shock measure is only available at the quarterly frequency, the surprise series is aggregated by summing across months.

C. Sensitivity Analysis

C.1. Instrument Construction

Event selection. A potential concern with our baseline instrument is that vessel collisions at maritime choke points may be more likely to occur during periods of heightened shipping activity, which could itself be correlated with broader macroeconomic conditions. To assess whether inclusion of collisions introduces any endogenous variation into our instrument, we re-estimate the responses using a version of the surprise series that excludes all collision-related events from the narrative dataset. As shown in Figure C.1, the impulse responses are virtually unchanged relative to the baseline, indicating that our results are not driven by the subset of events for which endogeneity concerns are most acute.

To additionally verify that no single event exerts undue influence on our estimates, we conduct a jackknife exercise in which we censor each observation in our surprise series to zero, one at a time, and re-estimate the baseline model for each resulting instrument. Figure C.2 shows the collection of responses using the modified instrument (in gray), along with the baseline response (in black). The results do not appear to be driven by any individual event, as the resulting responses lie safely within the 68 percent confidence interval of the baseline model.

Event source. We also assess whether the results are sensitive to the source of event information. Figure C.3 compares the baseline responses with responses based on instruments constructed separately from the Factiva–ACP dataset and the Lloyd's serious-casualty records. The impulse responses are similar across sources, indicating that the results are not driven by any single event source.

Predictability. Next, we examine whether the potential predictability of our events affects the results. A possible concern is that some disruptions may be at least partially anticipated. For instance, at the Panama Canal, draft restrictions imposed in response to adverse weather conditions may be anticipated as market participants may foresee the likelihood of further restrictions following an initial one. To address this concern, we construct a more restrictive instrument that retains only the first draft restriction in each sequence of related Panama Canal events, discarding subsequent restrictions that are arguably more predictable. As shown in Figure C.4, the impulse responses are very similar to the baseline, suggesting that anticipatory effects associated with repeated restrictions

Figure C.1: Sensitivity With Respect to Excluding Collisions

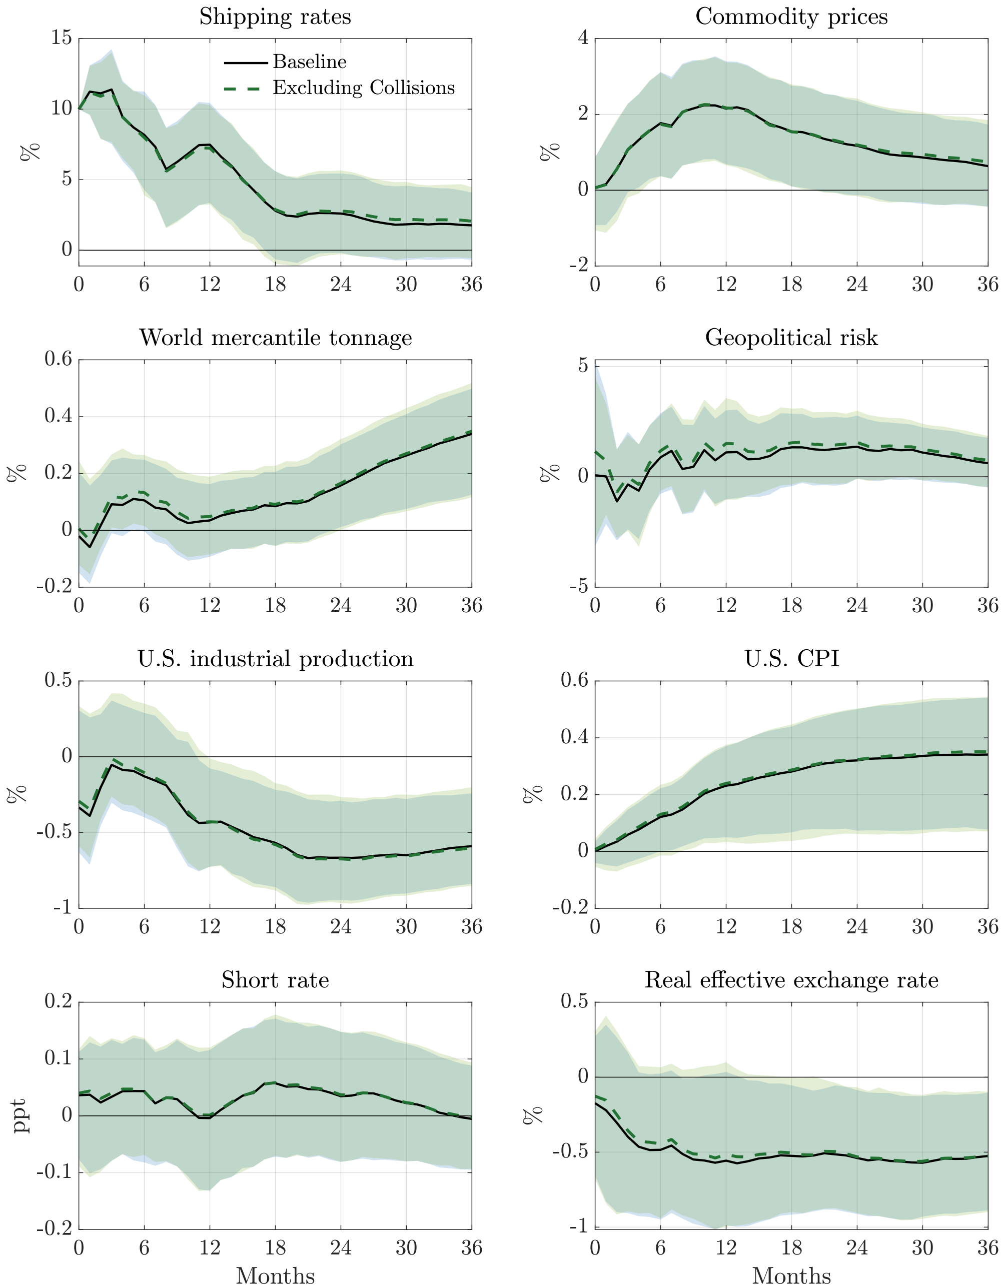

Notes: Impulse responses to a supply chain shock using an instrument that excludes collisions. The lines are point estimates and the shaded areas are 90 percent confidence bands.

Figure C.2: Jackknife Exercise

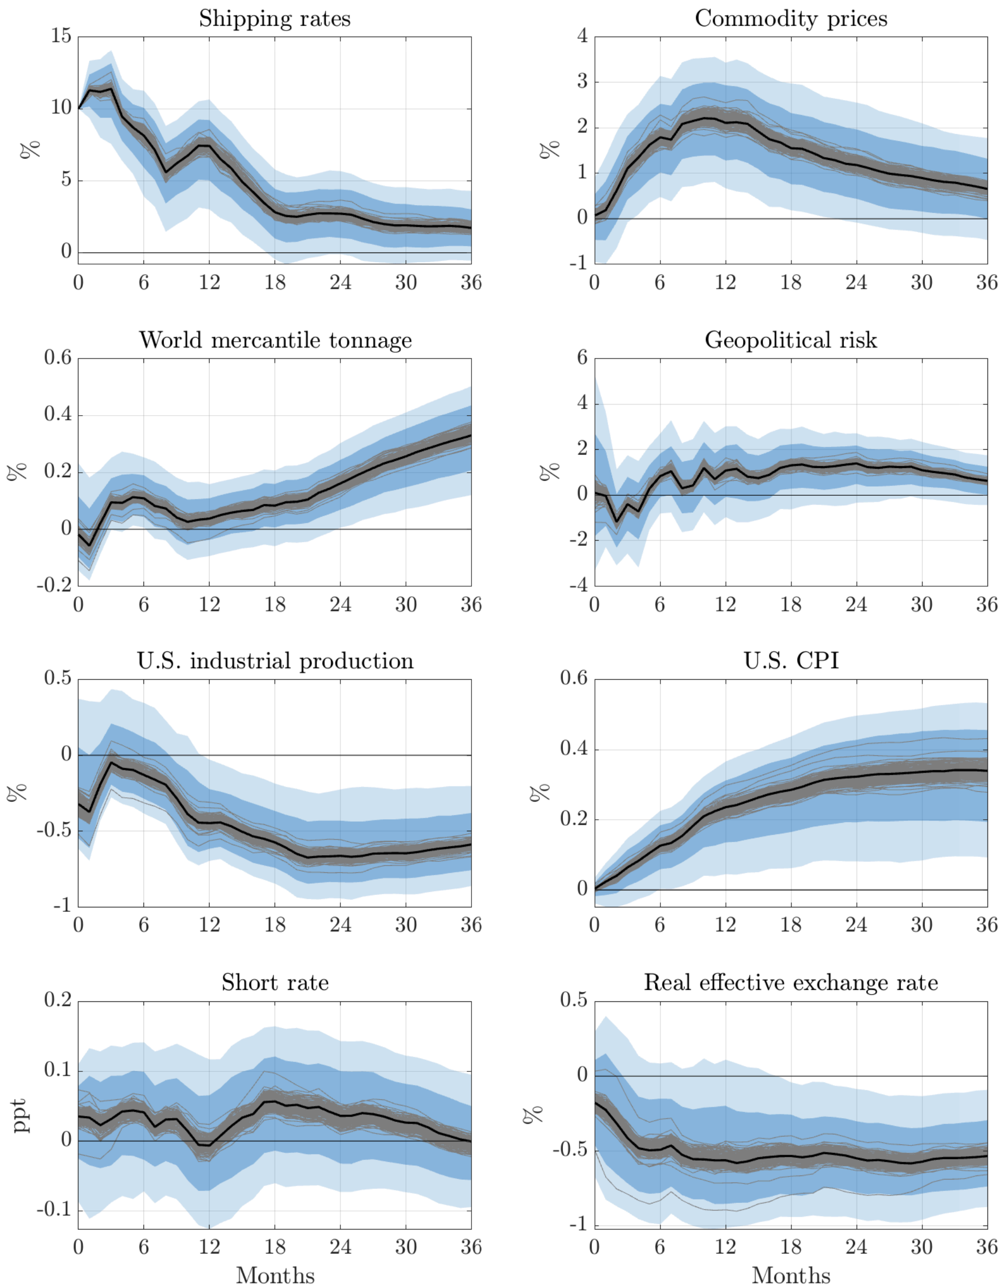

Notes: Impulse responses to a supply chain shock from the jackknife exercise. The lines are the point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands of the baseline model, respectively.

Figure C.3: Sensitivity With Respect to Event Source

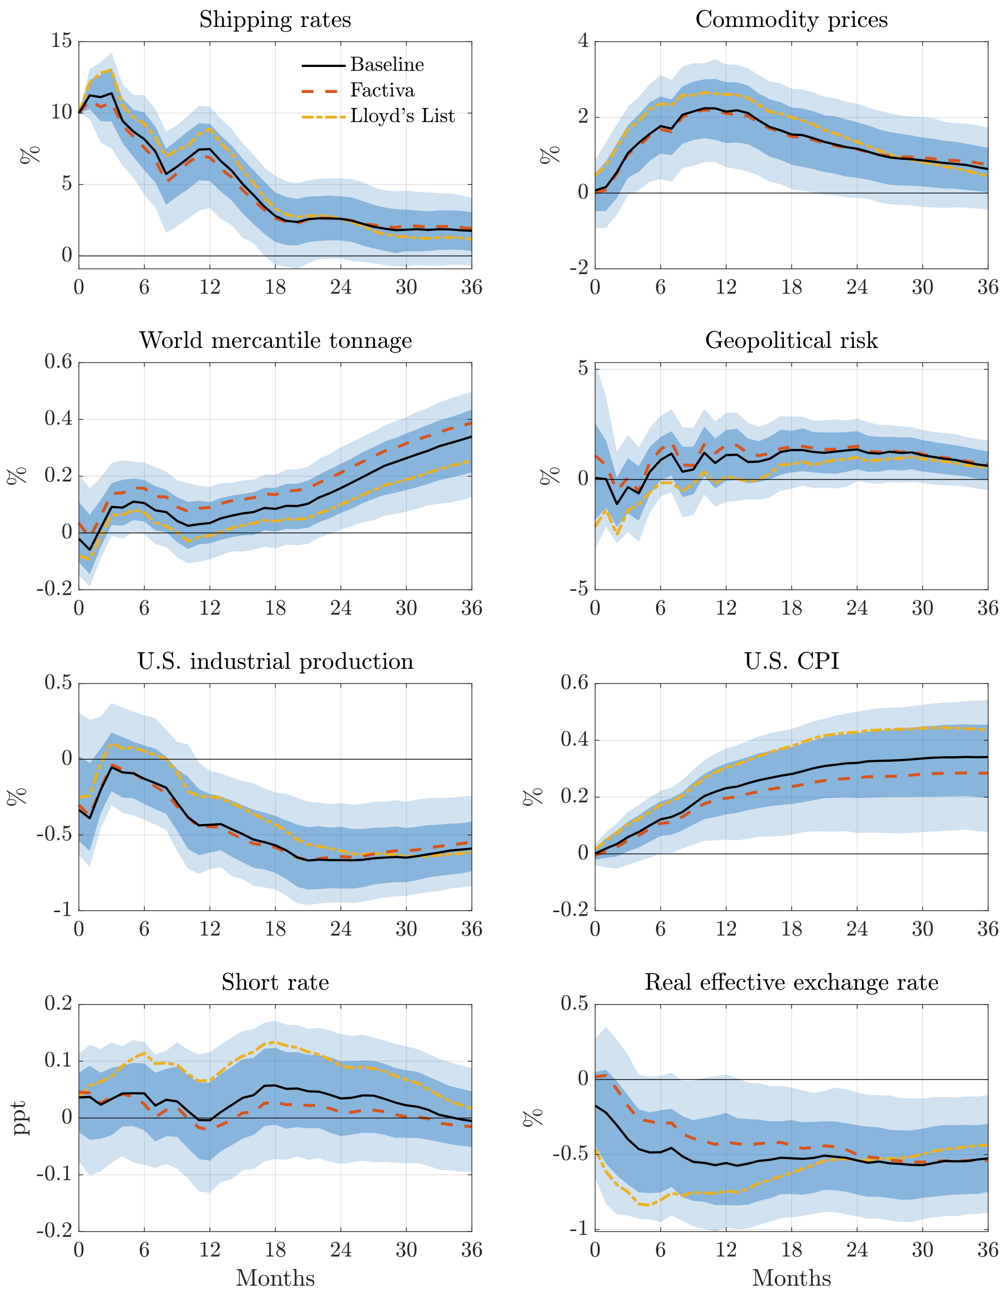

Notes: Impulse responses to a supply chain shock using an instrument that includes events in the Factiva–ACP dataset or in the Lloyd's dataset alone. The lines are the point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands of the baseline model, respectively.

do not materially alter our estimates.

As a complementary check, Figure C.5 shows the impulse responses using the raw surprise series, forgoing the attempts to account for predictability in Section 2.3. The results are virtually identical to the baseline, indicating that the predictability correction has little bearing on our estimates.

Event window. In our baseline, we measure shipping cost surprises using a one-day event window. To assess the sensitivity of our results to this choice, we re-estimate the model using surprises computed over a two-day window, which allows for the possibility that markets take longer to fully process certain disruptions. As shown in Figure C.6, the impulse responses obtained using the alternative window are essentially identical to the baseline, suggesting that the choice of window length does not materially affect our estimates.

Canals. A further concern is that disruptions in the Suez Canal may coincide with major geopolitical developments in the Middle East, raising the possibility that our instrument captures variation unrelated to supply chain conditions. While our baseline specification already controls for a news-based geopolitical risk index, we assess this more directly by re-estimating the model using instruments based on events in each canal separately. As shown in Figure C.7, the responses using only Panama Canal events are very similar to the baseline, alleviating concerns that geopolitical confounders are driving our results.

Negative surprises. A potential concern is that negative surprises in shipping costs may coincide with adverse macroeconomic news, confounding their interpretation as exogenous supply chain shocks. To mitigate this concern, we show that our baseline results are robust to keeping only positive shipping surprises, censoring negative surprises to zero. As shown in Figure C.8, the impulse responses are very similar to the baseline, though less precise as the instrument is constructed from a smaller set of events.

Aggregation. Figure C.9 shows the responses using instruments based on alternative methods of aggregating the daily surprise series to the monthly frequency. The cumulative-sum method assigns equal weight to all daily surprises within the month. The linear-weight method applies weights directly to daily surprises according to their position in the trading month. The triangular method follows the Gertler and Karadi (2015) aggregation approach, using changes in monthly averages of the cumulative daily surprise series. The impulse responses are similar across methods, suggesting that the results are

Figure C.4: Sensitivity With Respect to Event Selection

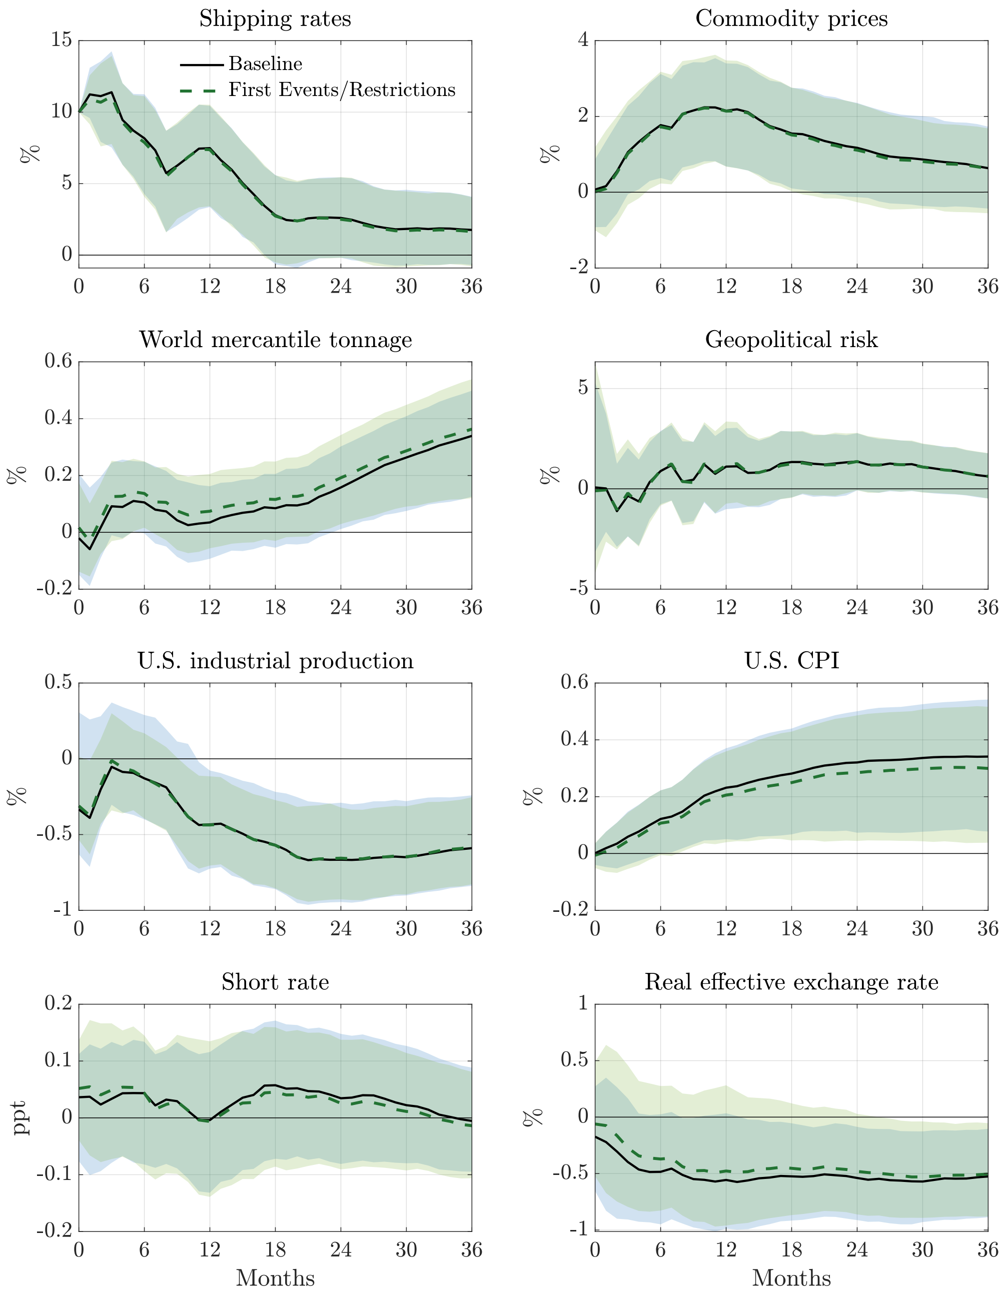

Notes: Impulse responses to a supply chain shock using an instrument based on a more restrictive set of events that are plausibly unanticipated. The lines are point estimates and the shaded areas are 90 percent confidence bands.

Figure C.5: Raw Surprises

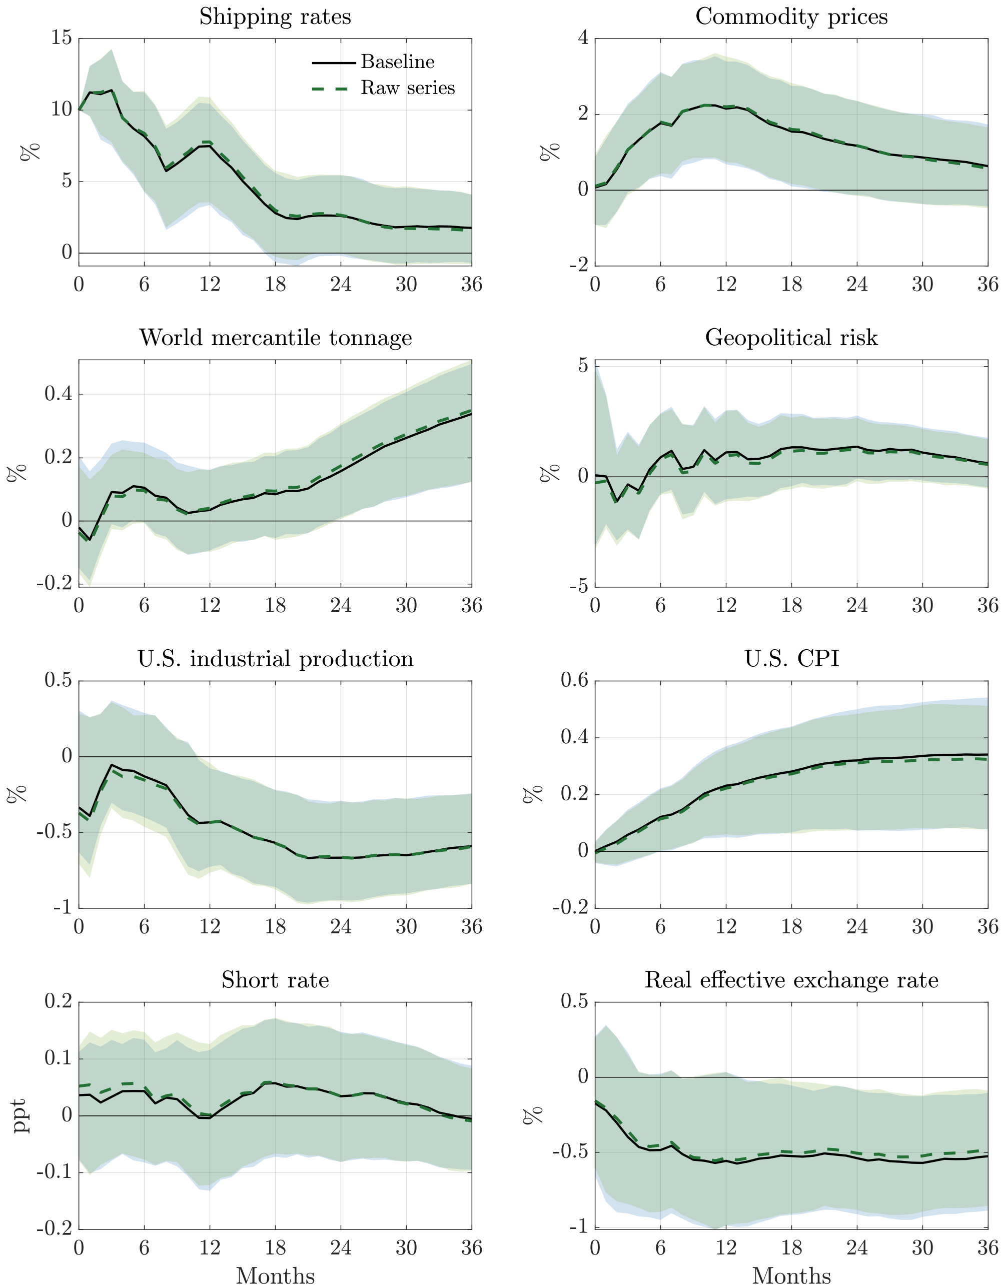

Notes: Impulse responses to a supply chain shock using the raw instrument. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

Figure C.6: Sensitivity With Respect to Event Window

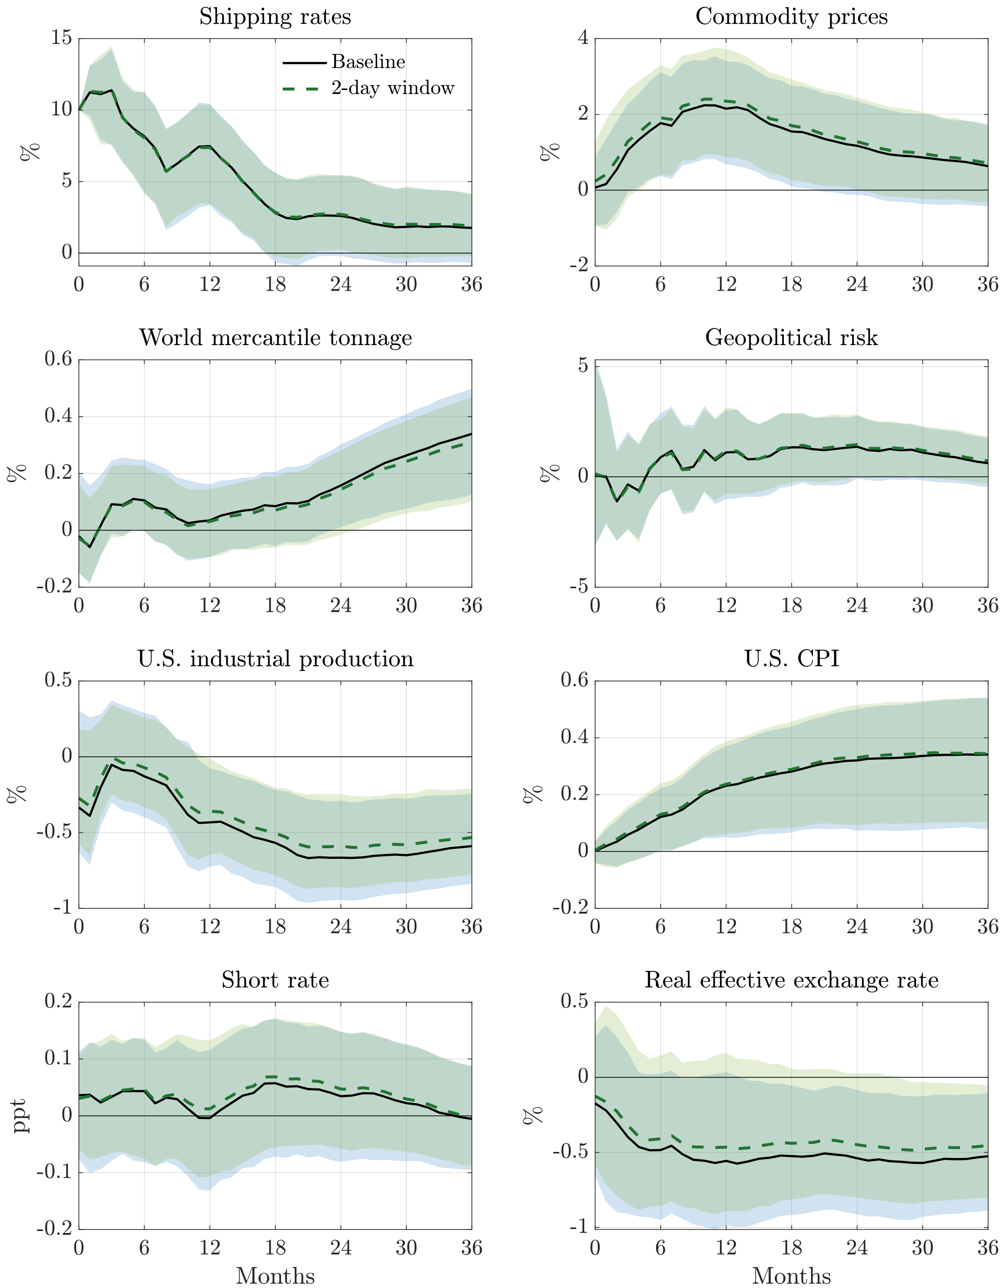

Notes: Impulse responses to a supply chain shock using an instrument with surprises computed over a 2-day event window. The lines are point estimates and the shaded areas are 90 percent confidence bands.

Figure C.7: Sensitivity With Respect to Canals

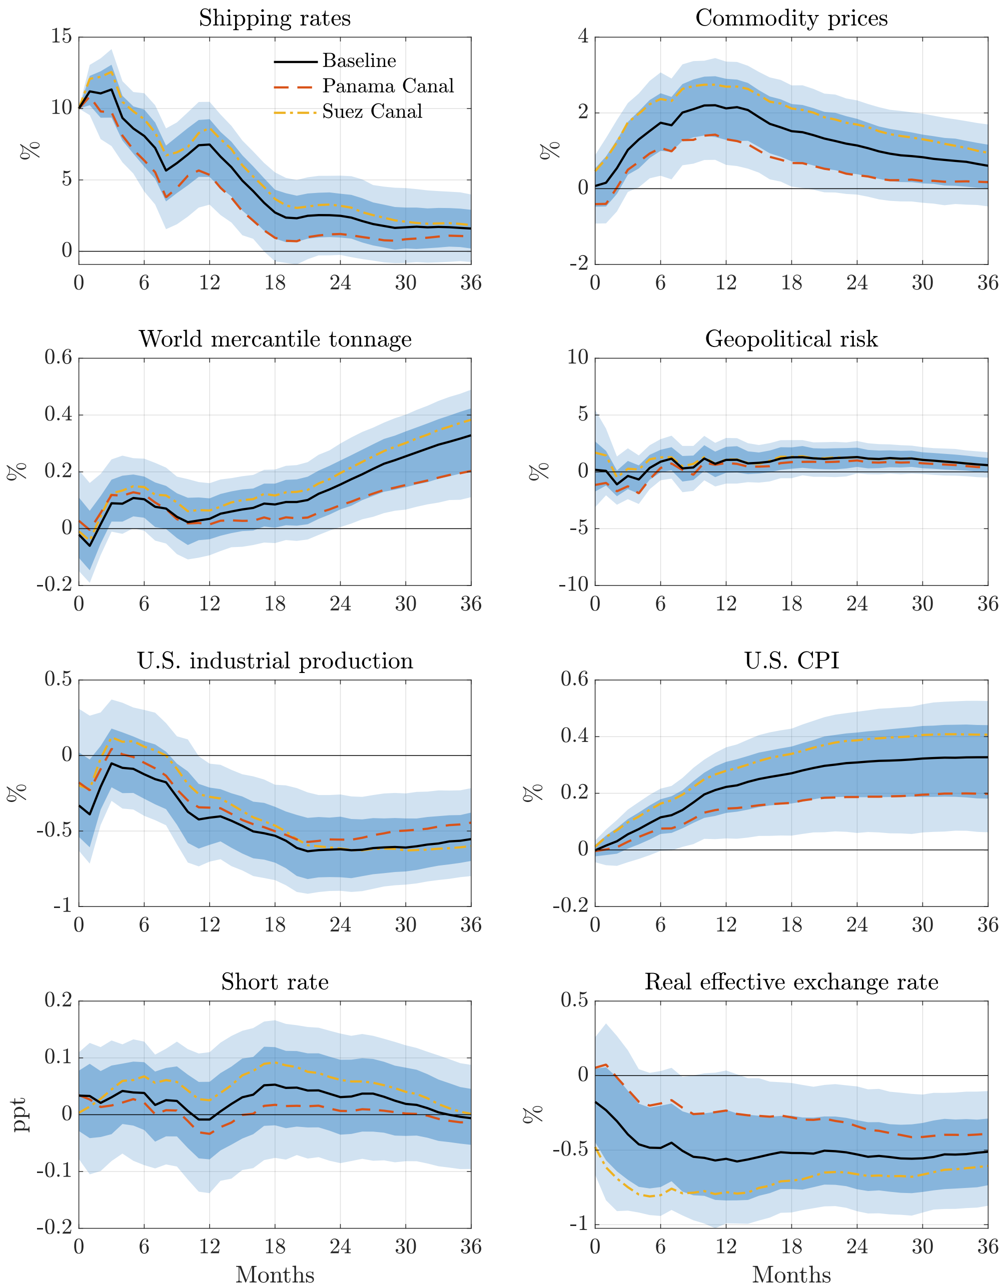

Notes: Impulse responses to a supply chain shock using an instrument that includes events in the Panama canal or in the Suez canal alone. The lines are the point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands of the baseline model, respectively.

Figure C.8: Sensitivity With Respect to Negative Surprises

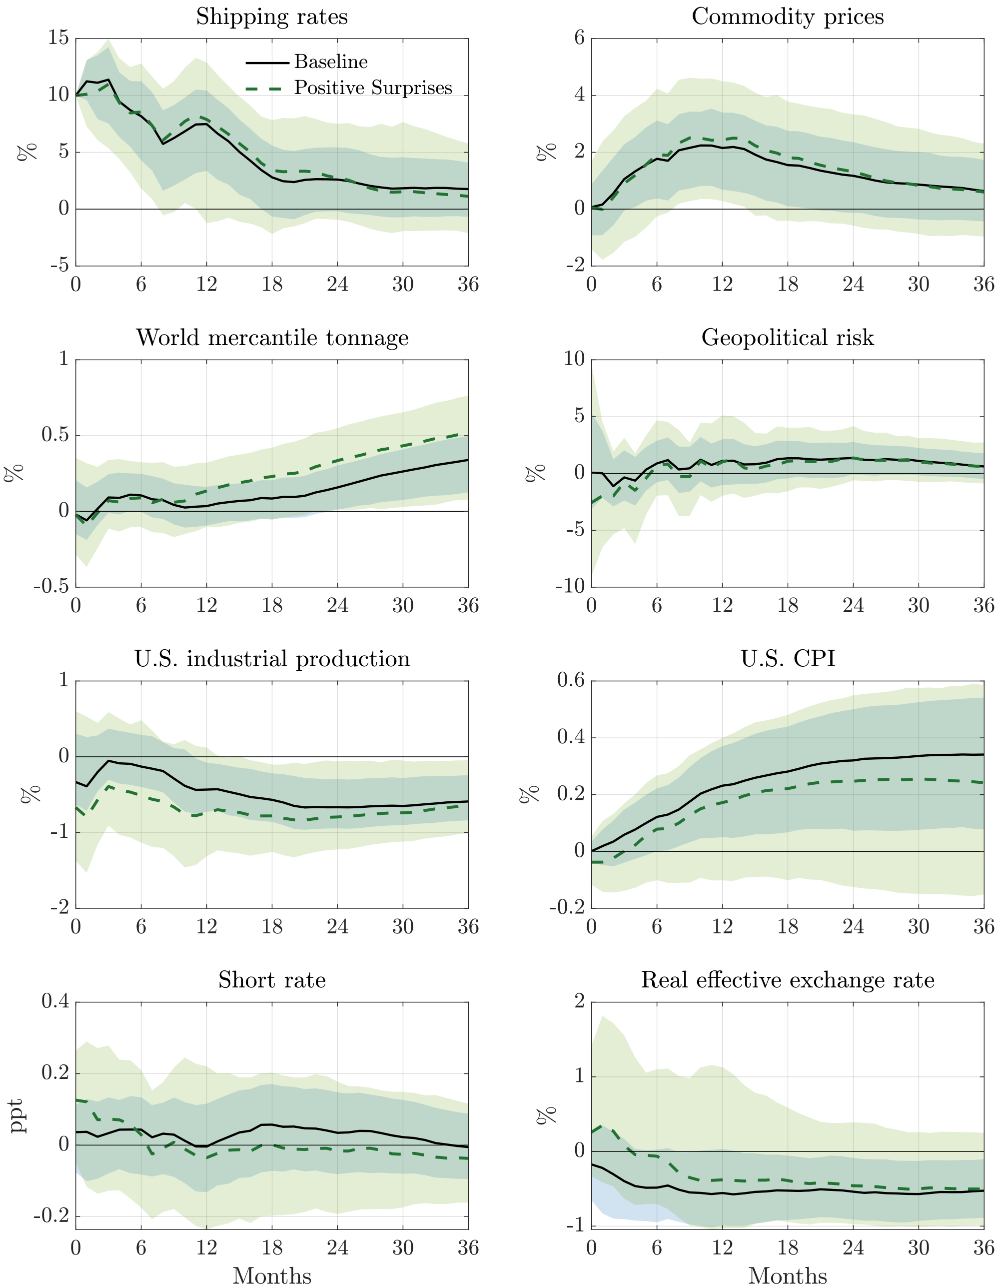

Notes: Impulse responses to a supply chain shock using an instrument that censors negative surprises to zero. The lines are point estimates and the shaded areas are 90 percent confidence bands.

not sensitive to the aggregation scheme.

C.2. Estimation Strategy

To mitigate concerns about invertibility, we perform two additional exercises. The first stays closer to the VAR framework but replaces the baseline proxy-VAR with two invertibility-robust alternatives: an internal instrument VAR and a generalized external-instrument VAR. The second estimates the dynamic effects directly using local projections with an external instrument.

Invertibility-robust internal- and external-instrument approaches. We implement two VAR-based procedures that are robust to non-invertibility. The first is the internal-instrument approach of Plagborg-Møller and Wolf (2019), which augments the VAR with the instrument ordered first and computes impulse responses to the first orthogonalized innovation.

The second is the generalized external-instrument VAR of Forni, Gambetti, and Ricco (2022), which allows the shock to be non-invertible but recoverable. When invertibility holds, this estimator reduces to the standard proxy-VAR. We first test directly for invertibility using their proposed diagnostic. Specifically, we regress the shipping-cost surprise z_t on the current reduced-form VAR innovations and on $r = 12$ monthly leads of those innovations, and then test the joint null that the coefficients on all future innovations are zero. In our baseline specification, the test fails to reject the null at conventional significance levels, providing no evidence against invertibility of the supply-chain shock.

Figure C.10 compares the responses from these estimation strategies with those from the baseline external-instrument VAR. The responses are broadly similar across the three specifications, suggesting that the results are not driven by the invertibility assumption.

Local projection-instrumental variable approach. As an invertibility-robust alternative, we estimate the impulse responses using a local projections–instrumental variable (LP-IV) approach à la Jordà, Schularick, and Taylor (2015) and Ramey and Zubairy (2018). Figure C.11 reports the resulting responses. The estimates are less precise than in the baseline VAR, as expected, but the qualitative pattern of the responses is broadly unchanged.

C.3. Sample and Specification Choices

Sample. We assess the sensitivity of our results to the sample period. The Covid-19 pandemic generated unprecedented disruptions to global supply chains that may be dif-

Figure C.9: Sensitivity With Respect to Aggregation Method for Daily Surprises

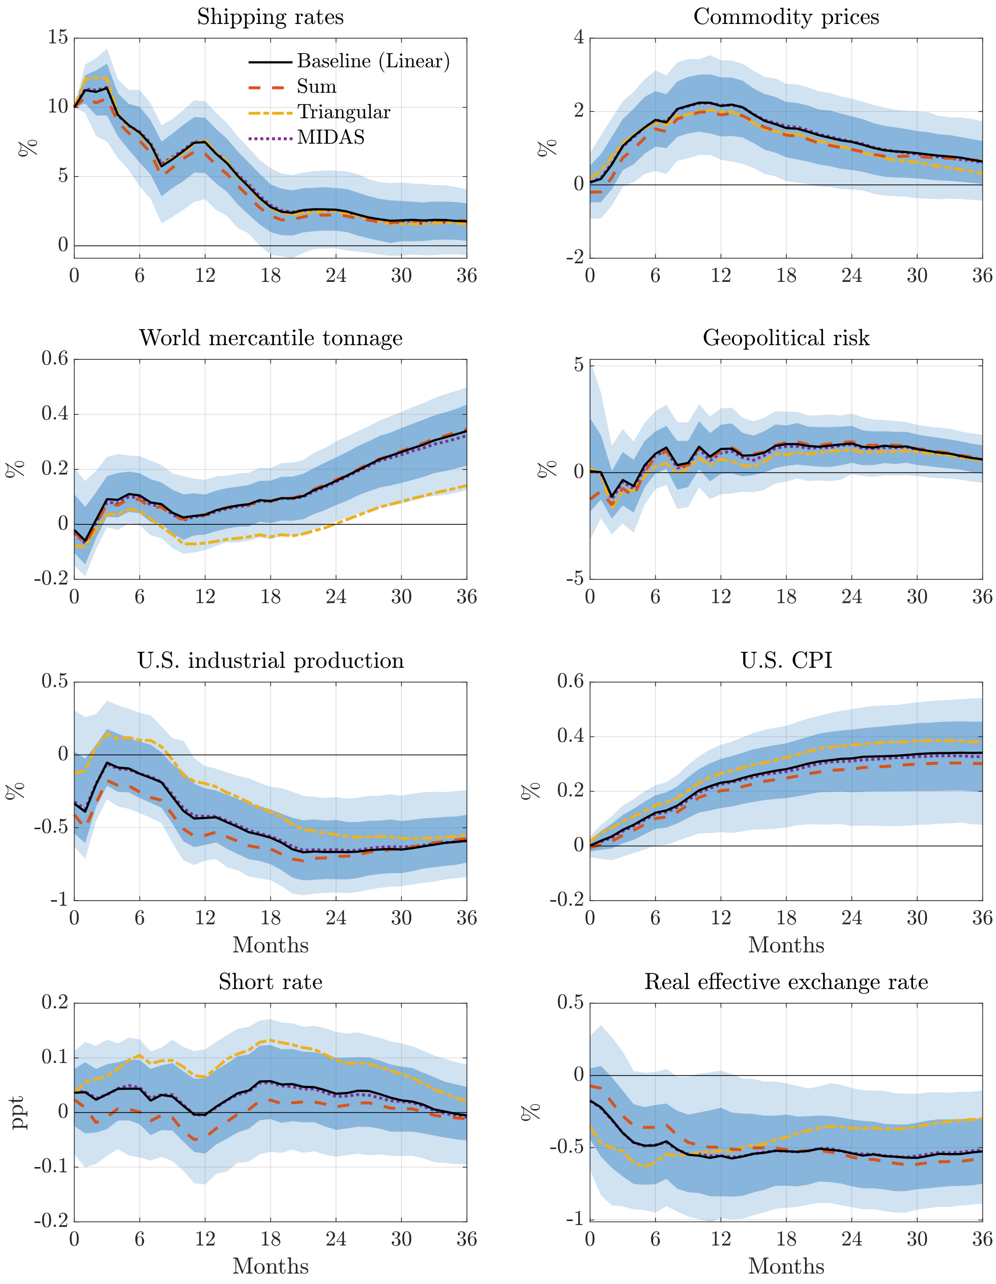

Notes: Impulse responses to a supply chain shock, using instruments based on different aggregation methods for the daily surprise series. The methods include cumulative sum (equal weights), linear weights, and triangular weights, all based on trading days. The lines are the point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands of the baseline model, respectively.

Figure C.10: Invertibility-Robust Internal and External Instruments Approach

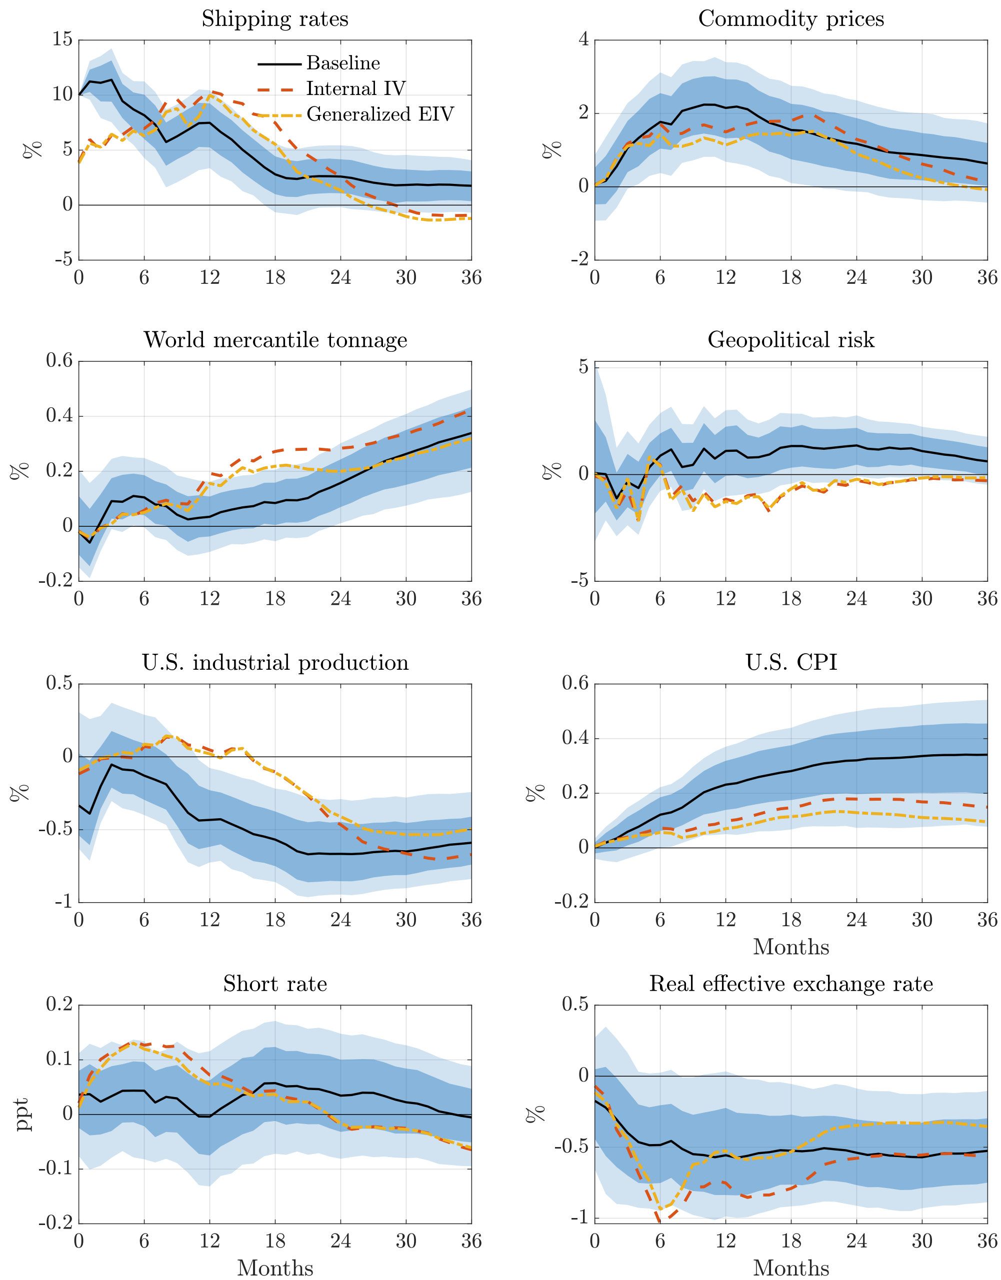

Notes: Impulse responses to a supply chain shock, estimated using the invertibility-robust internal and external instruments approach. The lines are the point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands of the baseline model, respectively.

Figure C.11: LP-IV Approach

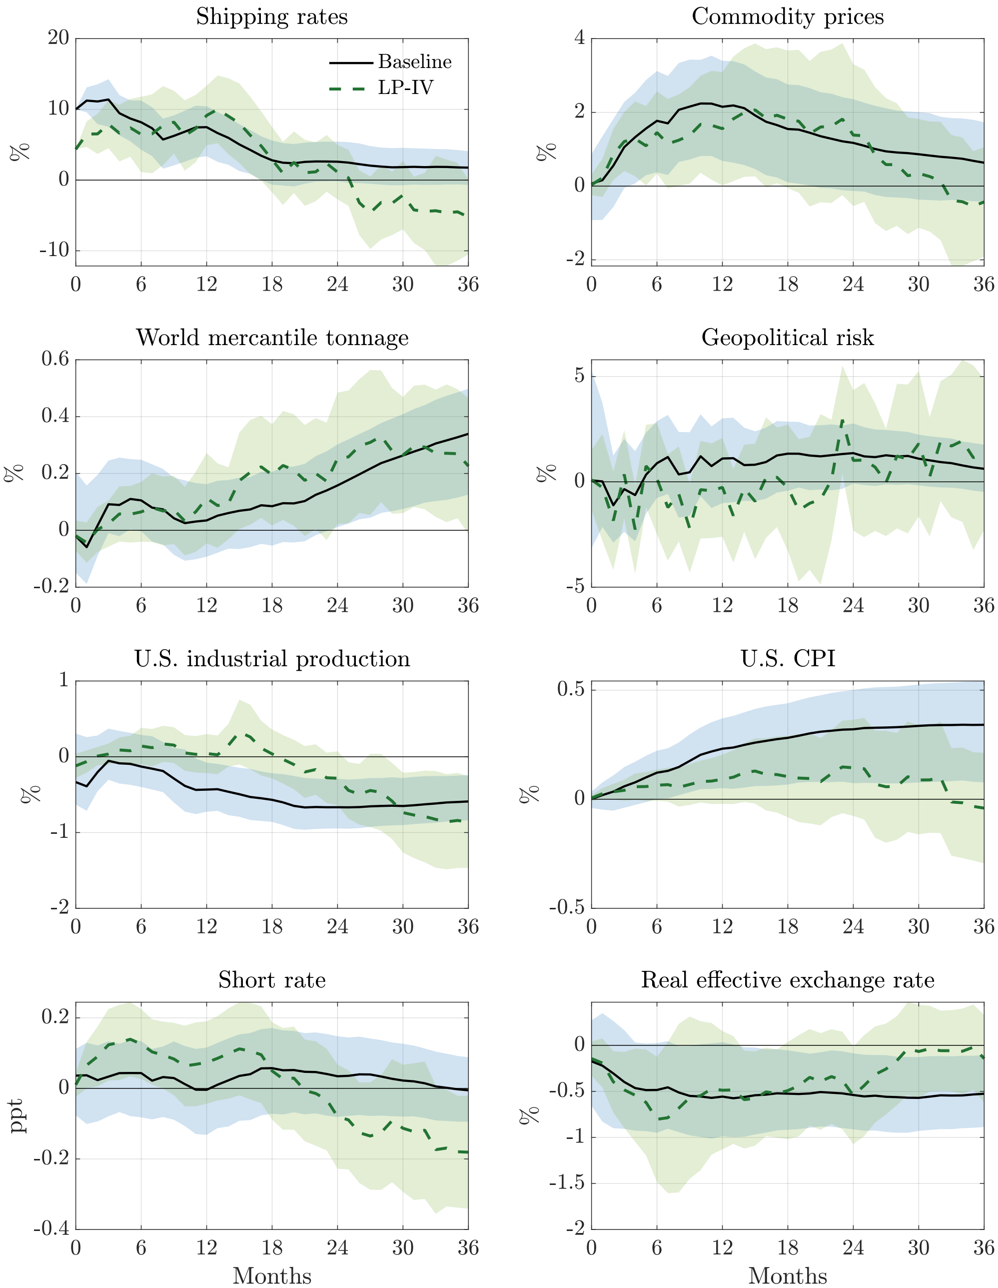

Notes: Impulse responses to a supply chain shock, estimated using the LP-IV approach. The lines are point estimates and the shaded areas are 90 percent confidence bands.

ficult to separate from the shipping disruptions captured by our instrument. Figure C.12 presents responses based on a sample excluding the Covid-19 period, and confirms that our baseline results are not driven by this episode.

Lag order. We assess the sensitivity of our results to the choice of lag order. Our baseline specification sets the lag order to 12, as is customary with monthly data. Figure C.13 shows the impulse responses under alternative lag orders of 6 and 18. The responses are similar across all three specifications, indicating that our results are not sensitive to the number of lags included in the VAR.

Deterministics. We also assess the sensitivity of our results to the choice of deterministic terms. Our baseline specification includes a linear trend. Figure C.14 shows the impulse responses when the linear trend is omitted and when we additionally control for the global financial crisis and the Covid-19 period using dummy variables. The responses are similar to the baseline across all specifications, indicating that our results are not driven by the particular choice of deterministic terms.

D. Additional Results

D.1. Aggregate Effects

Uncertainty. To sharpen the interpretation of the shock, we study how supply chain shocks influence a range of uncertainty and risk measures. Figure D.1 shows the responses of financial uncertainty, economic policy uncertainty, trade policy uncertainty, and crude oil volatility to a supply chain shock. None of these measures respond significantly, with the point estimates remaining close to zero and statistically insignificant throughout. This is reassuring, as it suggests that our identified shock represents a distinct source of supply chain disruption rather than a broader deterioration in economic or policy uncertainty, further supporting the validity of our identification strategy.

Industrial production. Figure D.2 shows the responses of industrial production categories by market group to a supply chain shock. Both final products and non-industrial supplies and materials decline significantly following the shock, consistent with the broad-based disruption to production activity that supply chain shocks generate. The responses are qualitatively similar across the two components, though the decline in intermediate inputs is notably larger. This pattern is intuitive—supply chain disruptions primarily im-

Figure C.12: Excluding the Covid-19 period

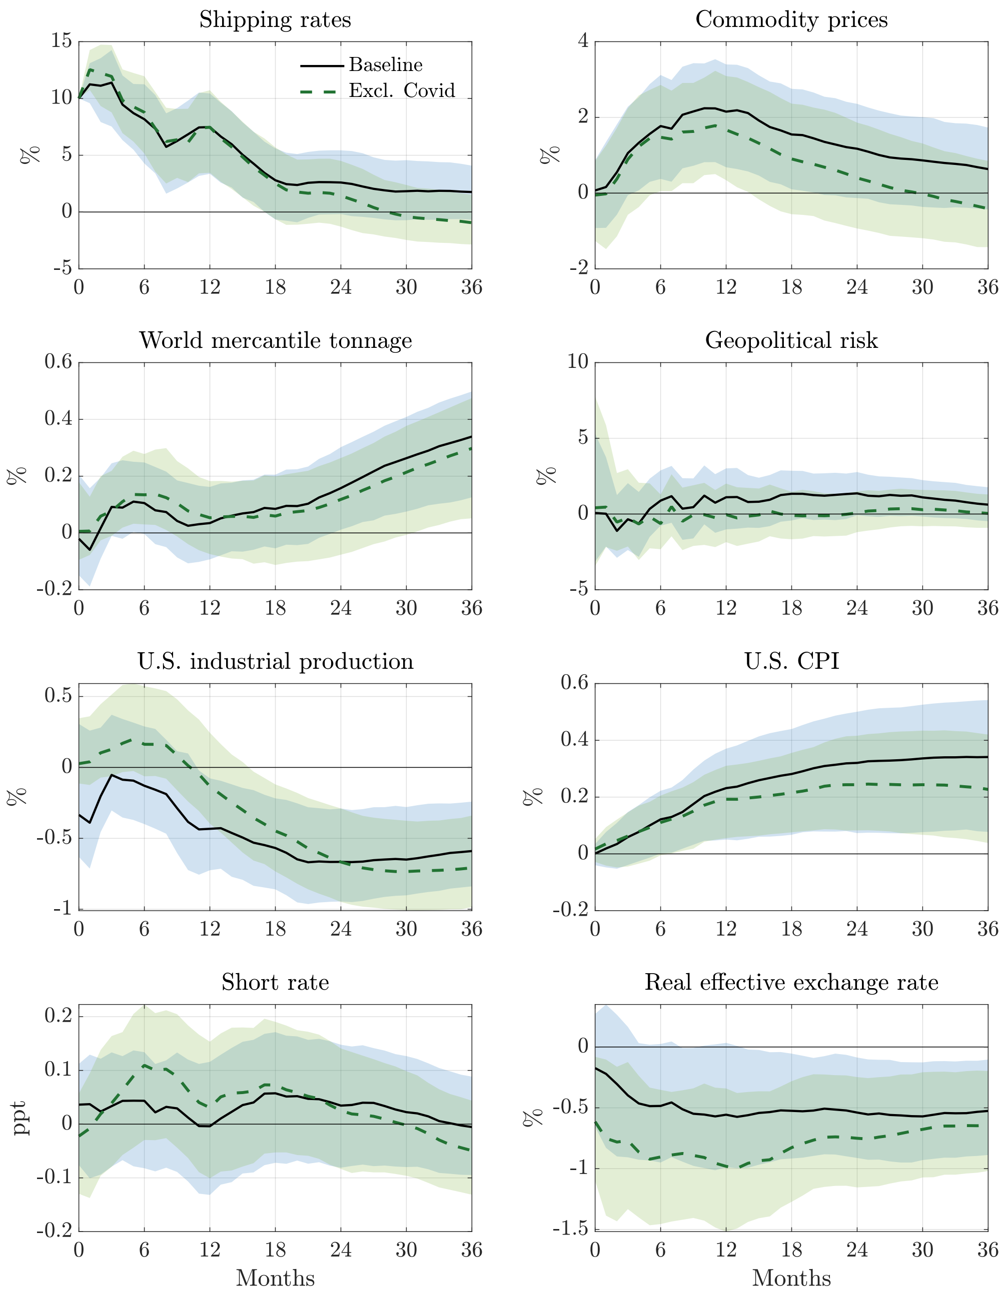

Notes: Impulse responses to a supply chain shock based on a sample excluding the Covid-19 period. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

Figure C.13: Sensitivity With Respect to Lag Order

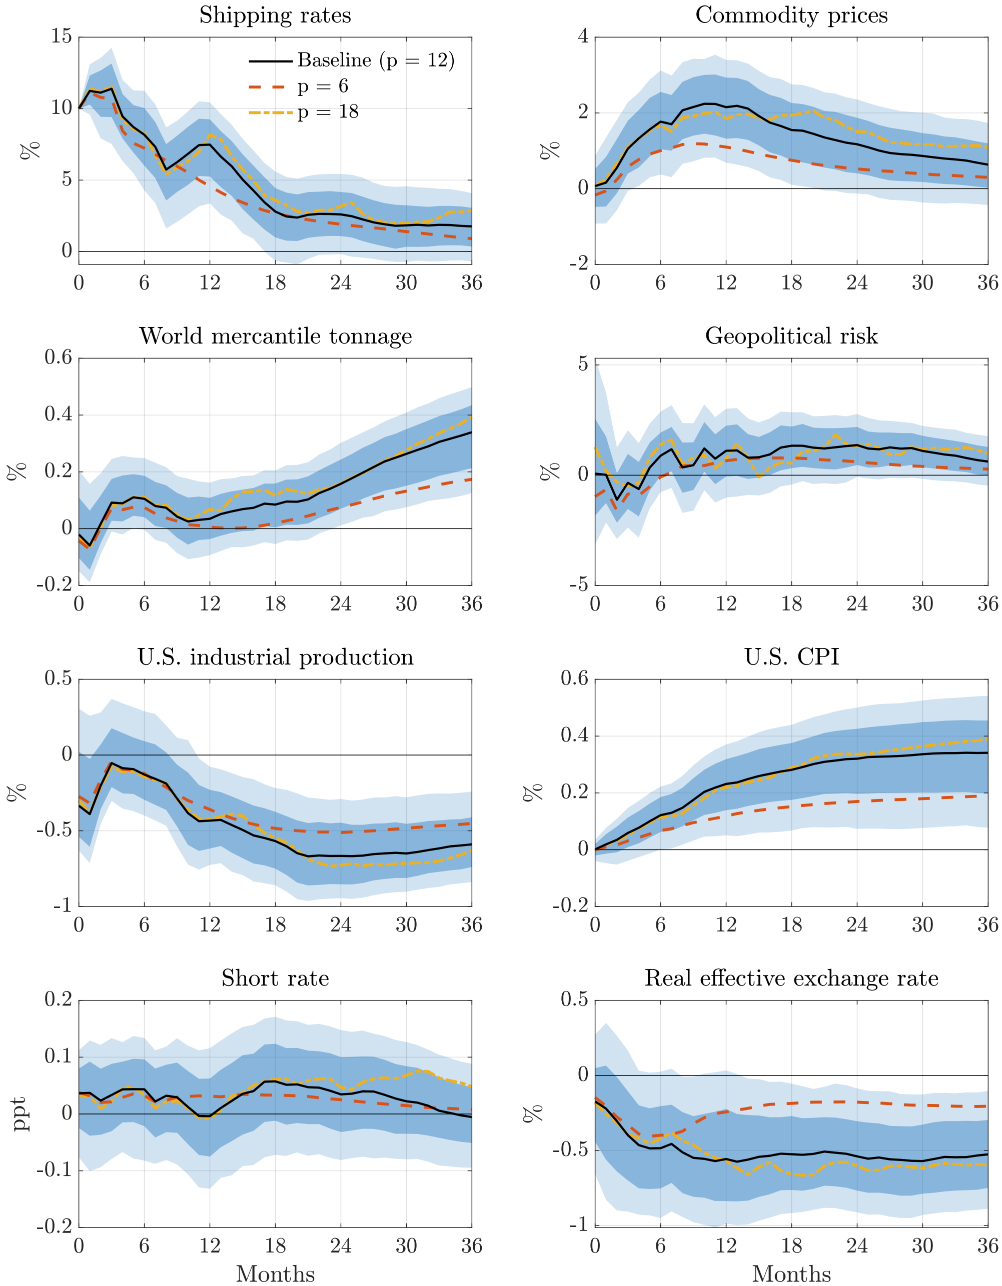

Notes: Impulse responses to a supply chain shock with varying lag order. The lines are the point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands of the baseline model, respectively.

Figure C.14: Sensitivity With Respect to Deterministics

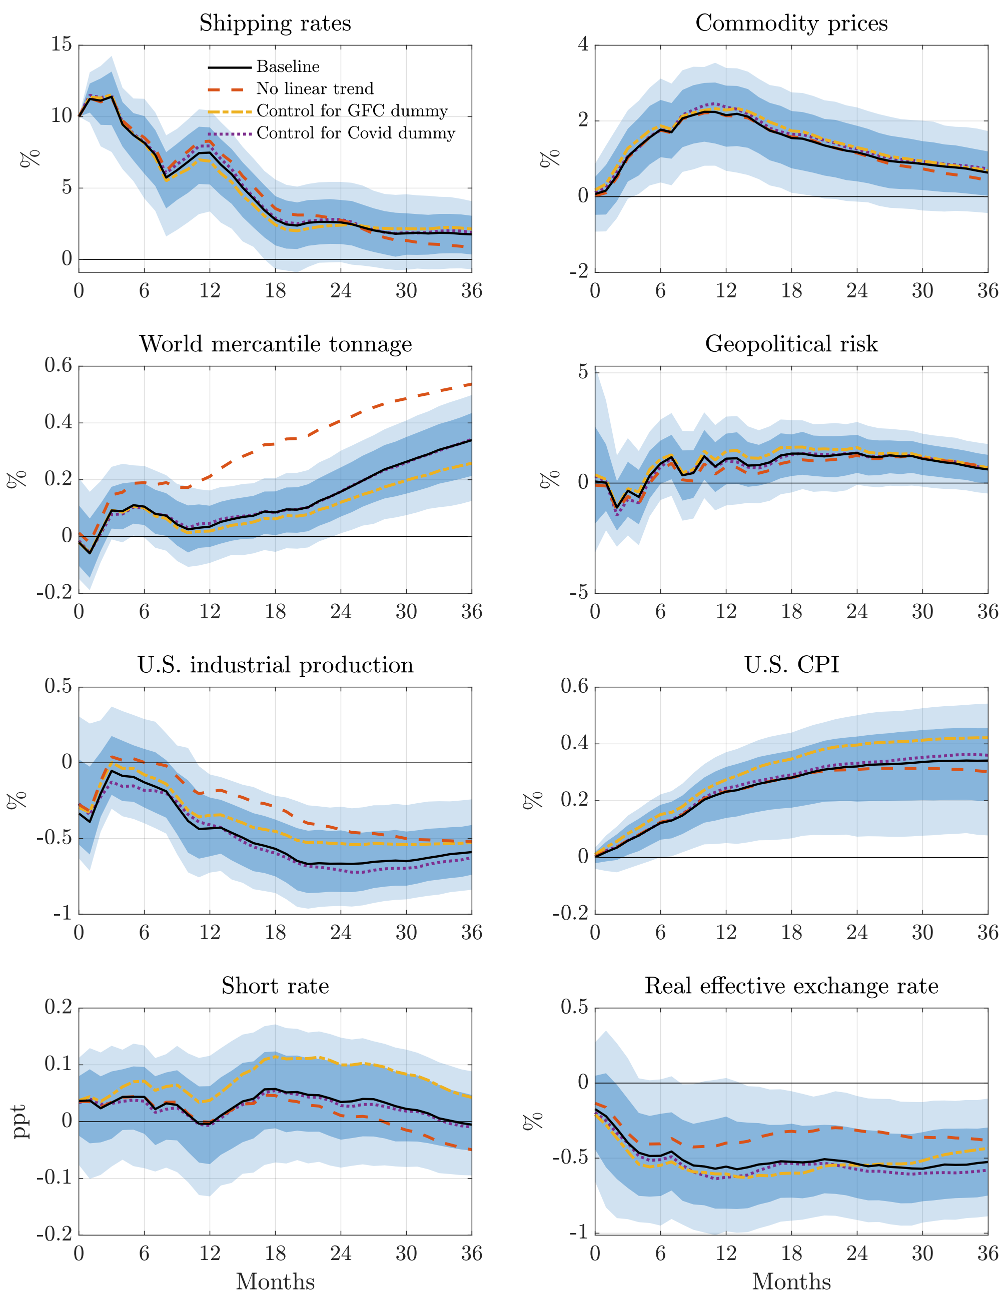

Notes: Impulse responses to a supply chain shock using a model that excludes the linear trend and controls for the global financial crisis (GFC) and Covid-19 using a dummy variable. The lines are the point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands of the baseline model, respectively.

Figure D.1: Impacts on Other Uncertainty and Risk Measures

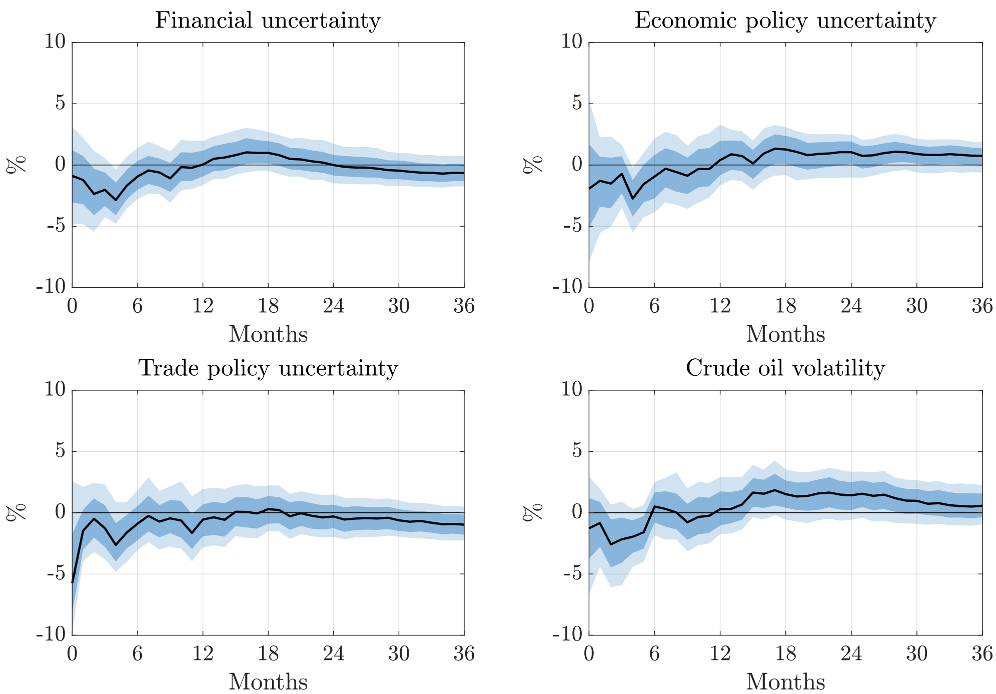

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The IRFs are obtained by using our baseline VAR and augmenting it by the variable of interest. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

pede the flow of intermediate inputs and raw materials, which bear the brunt of the shock before it propagates through to final goods production.

Figure D.2: Impact on Industrial Production Categories

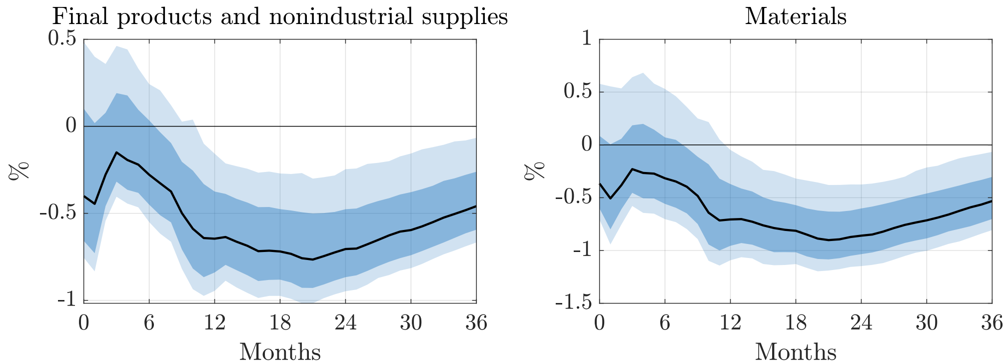

Notes: Impulse responses to a supply chain shock, normalized to increase real shipping costs by 10 percent on impact. The IRFs are obtained by using our baseline VAR and augmenting it by the variable of interest. The lines are point estimates and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

D.2. Sectoral Impacts and Substitution

One-step estimates of substitution elasticities. Section 5.2 estimates horizon-specific elasticities of substitution using a two-step procedure. The first step recovers, for each supplier-purchaser pair, the responses of relative input prices and expenditure shares to the identified supply-chain shock. The second step uses the cross-pair relationship between share responses and relative-price responses to recover the elasticity.

As a robustness exercise, we also estimate the same elasticity in a single-step local-projection instrumental-variables specification. This approach estimates the elasticity directly in the industry-pair panel, using the identified supply-chain shock interacted with predetermined supplier import exposure as an instrument for changes in relative input prices.

Specifically, for each horizon h, we estimate:

$$\ln s^M_{ij,t+h} - \ln s^M_{ij,t-1} = \alpha_i(h) + \lambda_t(h) + \theta(h)\Delta \ln \left(\frac{P_{jt}}{P^M_{it}}\right) + u_{ij,t+h}, \tag{3}$$

where i denotes the purchasing industry, j denotes the supplying industry, $s^M_{ij,t}$ is the share of purchaser i's intermediate-input expenditure allocated to supplier j, P_{jt} is the supplier gross-output price, and P^M_{it} is the purchaser intermediate-input price. As in the baseline exercise, the implied elasticity of substitution across intermediate suppliers is:

$$\eta(h) = 1 - \theta(h).$$

The relative input price term is potentially endogenous, since sectoral prices and sourcing shares may respond jointly to other shocks. We therefore instrument relative input prices with the interaction:

$$Z_{ij,t} = SCShock_t \times m_{j,t-1},$$

where $SCShock_t$ is the identified supply-chain shock and $m_{j,t-1}$ denotes the lagged share of supplier industry j's intermediate inputs that are imported.

Our preferred one-step specification includes purchasing-industry fixed effects, which absorb persistent differences across purchasers in production structure and input composition. We also estimate a version controlling for the supply chain shock, and an alternate version with year fixed effects. Since the aggregate supply-chain shock varies only over time, year fixed effects absorb the common aggregate component of the shock and any other annual macro disturbance. The identifying variation comes from differences in predetermined supplier exposure to the aggregate shock.

Figure D.3 displays the results. The fact that the estimates are similar with and without year fixed effects suggests that the one-step results are not driven by aggregate time-series confounders.

Figure D.3: Elasticity of Substitution Across Intermediate Inputs

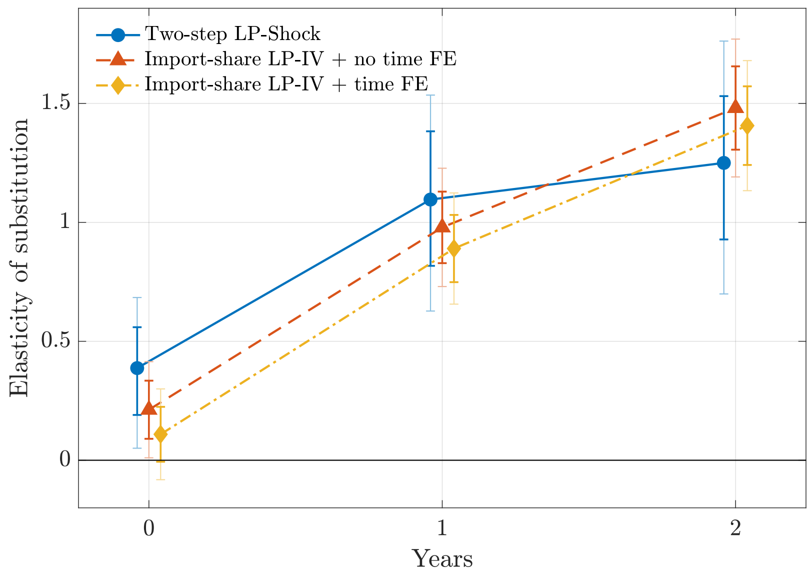

Notes: LP-IV estimates of the inner-nest elasticity of substitution, $\widehat{\eta}(h) = 1 - \widehat{\theta}(h)$. The LP-Shock specification (blue) is estimated using local projections (14) on the aggregated supply chain shock extracted from our baseline external instruments VAR. The import-share LP-IV specification is estimated using local projections (3), instrumenting relative input prices with our supply chain shock interacted with lagged supplier import shares. The "no time FE" specification (red) controls for the aggregate annual shock directly, while the "time FE" specification (yellow) absorbs year fixed effects. All specifications include purchaser-industry fixed effects. Vertical bars denote 68 and 90 percent confidence bands, respectively.

References Appendix

Atalay, Enghin (2017). "How important are sectoral shocks?" *American Economic Journal: Macroeconomics* 9.4, pp. 254–280.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis (2016). "Measuring economic policy uncertainty". *The Quarterly Journal of Economics* 131.4, pp. 1593–1636.

Barsky, Robert B. and Eric R. Sims (2011). "News shocks and business cycles". *Journal of Monetary Economics* 58.3, pp. 273–289.

Bassett, William F., Mary Beth Chosak, John C. Driscoll, and Egon Zakrajšek (2014). "Changes in bank lending standards and the macroeconomy". *Journal of Monetary Economics* 62, pp. 23–40.

Basu, Susanto, John G. Fernald, and Miles S. Kimball (2006). "Are technology improvements contractionary?" *American Economic Review* 96.5, pp. 1418–1448.

Bauer, Michael D. and Eric T. Swanson (2023). "A reassessment of monetary policy surprises and high-frequency identification". *NBER Macroeconomics Annual* 37.1, pp. 87–155.

Baumeister, Christiane and James D. Hamilton (2019). "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks". *American Economic Review* 109.5, pp. 1873–1910.

Beaudry, Paul and Franck Portier (2014). "News-driven business cycles: Insights and challenges". *Journal of Economic Literature* 52.4, pp. 993–1074.

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

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello (2019). "Oil price elasticities and oil price fluctuations". *Journal of Monetary Economics* 103, pp. 1–20.

Caldara, Dario and Matteo Iacoviello (2022). "Measuring Geopolitical Risk". *American Economic Review* 112.4, pp. 1194–1225.

Caldara, Dario, Matteo Iacoviello, Patrick Molligo, Andrea Prestipino, and Andrea Raffo (2020). "The economic effects of trade policy uncertainty". *Journal of Monetary Economics* 109, pp. 38–59.

Caldara, Dario, Matteo Iacoviello, and David Yu (2024). "Measuring Shortages Since 1900". *Working Paper*.

Fisher, Jonas D.M. and Ryan Peters (2010). "Using stock returns to identify government spending shocks". *The Economic Journal* 120.544, pp. 414–436.

Forni, Mario, Luca Gambetti, and Giovanni Ricco (2022). "External instrument svar analysis for noninvertible shocks".

Gertler, Mark and Peter Karadi (2015). "Monetary policy surprises, credit costs, and economic activity". *American Economic Journal: Macroeconomics* 7.1, pp. 44–76.

Gilchrist, Simon and Egon Zakrajšek (2012). "Credit spreads and business cycle fluctuations". *American Economic Review* 102.4, pp. 1692–1720.

Hamilton, James D. (2021). "Measuring global economic activity". *Journal of Applied Econometrics* 36.3, pp. 293–303.

Horowitz, Karen J. and Mark A. Planting (2009). "Concepts and Methods of the U.S. Input-Output Accounts". *Bureau of Economic Analysis*.

Jacks, David S. and Martin Stuermer (2021). "Dry bulk shipping and the evolution of maritime transport costs, 1850–2020". *Australian Economic History Review* 61.2, pp. 204–227.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor (2015). "Betting the house". *Journal of International Economics* 96, S2–S18.

Känzig, Diego R. (2021). "The macroeconomic effects of oil supply news: Evidence from OPEC announcements". *American Economic Review* 111.4, pp. 1092–1125.

Kurmann, André and Christopher Otrok (2013). "News shocks and the slope of the term structure of interest rates". *American Economic Review* 103.6, pp. 2612–2632.

Miranda-Agrippino, Silvia and Giovanni Ricco (2021). "The Transmission of Monetary Policy Shocks". *American Economic Journal: Macroeconomics* 13.3, pp. 74–107.

Plagborg-Møller, Mikkel and Christian K. Wolf (2019). "Local projections and VARs estimate the same impulse responses".

Ramey, Valerie A. (2011). "Identifying government spending shocks: It's all in the timing". *The Quarterly Journal of Economics* 126.1, pp. 1–50.

Ramey, Valerie A. and Sarah Zubairy (2018). "Government spending multipliers in good times and in bad: evidence from US historical data". *Journal of Political Economy* 126.2, pp. 850–901.

Romer, Christina D. and David H. Romer (2004). "A new measure of monetary shocks: Derivation and implications". *American Economic Review* 94.4, pp. 1055–1084.

— (2010). "The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks". *American Economic Review* 100.3, pp. 763–801.

Smets, Frank and Rafael Wouters (2007). "Shocks and frictions in US business cycles: A Bayesian DSGE approach". *American Economic Review* 97.3, pp. 586–606.