

The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory*

Xiwen Bai[†]

Jesús Fernández-Villaverde[‡]

Yiliang Li[§]

Francesco Zanetti[¶]

December 5, 2025

Abstract

We study the causal effects of global supply chain disruptions by constructing a new index of real-time port congestion using Automatic Identification System data from container ships and a spatial clustering algorithm. We develop a model with search frictions between producers and retailers that links upstream production slack to downstream supply shortages and captures output and price responses to supply chain shocks. The co-movements of output, prices, spare capacity, and market tightness provide novel identification restrictions. We find that demand and supply shocks drove U.S. disinflation in 2020, while the inflation surge in 2021 was driven mainly by supply chain shocks.

JEL Classification: C32, E31, E32, J64.

Keywords: Supply chain disruption, satellite data, spatial clustering method, search-and-matching in the goods market, SVAR.

*We thank the editor, Yuriy Gorodnichenko, and four anonymous referees for their valuable suggestions. We are also grateful to Klaus Adam, Hassan Afrouzi, George Alessandria, Fernando Álvarez, Yan Bai, Gianluca Benigno, Hilde Bjørnland, Dennis Bonam, Bjoern Bruegemann, Vasco Carvalho, Xiaomin Cui, Elena María Díaz, Julian di Giovanni, Gauti Eggertsson, Mariassunta Giannetti, Mishel Ghassibe, Sebastian Heise, Bo Hu, Zhen Huo, Zuhuan Jin, Callum Jones, Şebnem Kalemlı-Özcan, Enisse Kharroubi, Nobuhiro Kiyotaki, Andras Komaromi, Laura Lebastard, Andrei Levchenko, Ernest Liu, Marco Lombardi, Benjamin Moll, Emi Nakamura, George Nikolakoudis, Theodore Papageorgiou, Ekaterina Peneva, Justin Pierce, Giorgio Primiceri, Xincheng Qiu, Omar Rachedi, Stephen Redding, Ricardo Reis, Esteban Rossi-Hansberg, Kathryn Russ, Thomas Sargent, Karthik Sastry, Adam Shapiro, Hyun Song Shin, Frank Smets, Vladimir Smirnyagin, Bo Sun, Aleh Tsyvinski, Harald Uhlig, Cliff Winston, Christian Wolf, Woan Foong Wong, Le Xu, Penghui Yin, Changhua Yu, Li Yu, Yang Yu, Bianca Zanetti, Charles Zhang, Xin Zhang, Yuan Zi, and participants at numerous conferences and seminars for their comments and suggestions. Zhongjun Ma provided excellent research assistance. Francesco Zanetti gratefully acknowledges financial support from the British Academy.

† Bai: Tsinghua University, China. xiwenbai@mail.tsinghua.edu.cn. ‡ Fernández-Villaverde: University of Pennsylvania, U.S. jesusfv@econ.upenn.edu. § Li: University of International Business and Economics, China. yiliang_li@uibe.edu.cn. ¶ Zanetti: University of Oxford, U.K. francesco.zanetti@economics.ox.ac.uk.

1. Introduction

The world economy is organized around an intricate global supply chain. Any sudden and large shock to this supply chain, such as those triggered by war, the COVID-19 pandemic, or the Red Sea crisis, can have significant consequences for output, inflation, and economic slack.

Measuring the causal effects of a global supply chain shock is challenging for two reasons. First, researchers need an accurate gauge of supply chain conditions, yet existing indices often rely on shipping prices or survey data from the Purchasing Managers' Index (PMI). Although informative, these measures may be skewed by irrelevant factors and large measurement errors. Ideal measurement requires precise, real-time data tracking of global flows of tradable goods.

Second, researchers need a theoretical framework that can deliver credible identification assumptions for causal analysis. Global supply chain conditions respond to aggregate demand, aggregate supply, and supply chain shocks, which can only be disentangled through theoretically derived identification restrictions. However, no standard model captures the joint presence of upstream economic slack and downstream supply scarcity, a combination that, as we argue later, is crucial for distinguishing supply chain disturbances from other macroeconomic shocks.

Our paper tackles these challenges by developing (i) a new index that measures real-time container ship congestion at major ports worldwide using high-frequency satellite data, providing a timely and accurate indicator of the state of the global supply chain, and (ii) a novel theoretical framework that captures the coexistence of upstream slack and downstream shortages and examines their implications for output, prices, market tightness, and spare productive capacity during supply chain disruptions. Using identification assumptions grounded in this theory, together with a Bayesian structural vector autoregression (SVAR), we disentangle the shocks driving our index and quantify the dynamic causal impact of supply chain shocks on aggregate outcomes.

The importance of addressing points (i) and (ii) lies in the likelihood that the world economy may again face major supply chain disruptions —whether from wars, geostrategic shifts, blockades, sanctions, or another pandemic. Far from being a postmortem of the COVID-19 pandemic, our analysis distills important lessons for the future.

Measuring the state of the global supply chain. We assess the health of the global supply chain by examining congestion at container ports worldwide, a well-recognized metric in maritime economics. As early as 2006, the [Transportation Research Board Executive Committee](#)

identified congestion as a key issue for transportation and logistics, a view later reinforced by [Fan et al. \(2012\)](#) and [Brancaccio et al. \(2024\)](#), who documented its impact on the efficiency and reliability of global supply chains.

Container shipments are central to global trade, with roughly 60% of the total value of seaborne trade passing through container ports ([UNCTAD, 2019](#)). This heavy reliance on containerized transport implies that even small increases in port congestion can disrupt the supply–demand balance for tradable goods and strain the global supply chain.

Our port congestion analysis relies on real-time, high-frequency satellite data from the Automatic Identification System (AIS), mandated by the International Maritime Organization (IMO). These data allow us to track container ships with virtually no measurement error between 2016 and 2025. We measure congestion in individual ports using a machine-learning-based spatial clustering algorithm that leverages ship positions, speeds, and headings, and we aggregate these port-level measures to construct the first high-frequency Average Congestion Rate (ACR) index.

Our index shows that COVID-19-related port congestion began rising in the second half of 2020 and remained elevated until mid-2022. The share of delayed container ships increased from 28% to 37%, while average delay durations rose from 6 to 14 hours. Given that almost 80% of global trade is shipped indirectly, with an average of five port stops ([Ganapati et al., 2024](#)), these seemingly small delays imply substantial disruptions in container flows.

A model of the global supply chain. Next, we develop a model to capture imbalances between goods supply and demand during supply chain disturbances. The model centers on search and matching frictions between geographically separated producers and retailers, with producers incurring transportation costs to ship goods.

Our framework is inspired by the disequilibrium literature of the 1970s (e.g., [Barro and Grossman 1971](#)) but is recast within the microfounded approach of [Michaillat and Saez \(2015, 2022\)](#) and [Ghassibe and Zanetti \(2022\)](#). By distinguishing between producers and retailers and incorporating transportation costs, the model generates spare productive capacity in the upstream producer–retailer market and supply shortages in the downstream retailer–household market. Search frictions introduce trading externalities that constrain the allocative role of prices. In this setting, trading depends on the relative numbers of retailers and producers, in addition to standard price adjustments, a crucial mechanism during supply chain disruptions.

We model a supply chain disturbance as an unexpected increase in transportation costs.

Such cost spikes —driven by port congestion surcharges, shipping shortages, or pandemic-era price increases (Alessandria et al., 2023; Dunn and Leibovici, 2023)— reduce the expected total surplus from potential producer–retailer matches. Because the maximum transportation cost that a match can bear adjusts sluggishly under fixed service contracts, the upstream producer–retailer market slackens, generating spare productive capacity while lowering the supply of goods and raising prices in the downstream retailer–household market.

Identification. Identification in our estimation below relies on two sources: the restrictions derived from our model and our domain knowledge of the containerized shipping industry.

The restrictions implied by our model highlight that aggregate responses to supply chain shocks differ clearly from those to conventional demand or supply shocks. Specifically, supply chain shocks generate negative co-movements between output and prices, similar to traditional supply shocks. However, unlike standard supply shocks, supply chain disturbances also lead to both an increase in *global* spare capacity (as reduced shipments constrain goods flows to retailers without affecting productive capacity) and a decrease in *upstream* market tightness (as retailers are discouraged from trade while spare productive capacity rises). This combination of higher spare capacity, lower market tightness, rising prices, and declining output enables us to uniquely identify supply chain disturbances.

Domain knowledge tells us that port congestion is unaffected by aggregate demand or productive capacity shocks *within the first month* after such shocks occur. This is due to two main factors: (i) container ship schedules and arrivals are determined ex-ante and adjust to changes in demand or capacity only after significant delays, driven by penalties and high switching costs, and (ii) travel times between ports often exceed a month, making it impossible for container ships already en route to respond to demand or capacity changes in under 30 days.¹

We exploit this delayed response by imposing that the condition that the ACR index remains unresponsive to aggregate demand or productive capacity shocks during the first month post-impact, while allowing it to respond freely thereafter. Importantly, as we elaborate below, our ACR index, combined with this domain knowledge identification, delivers substantial gains in separating supply chain disturbances from other shocks. These gains arise both in the magnitude of price responses to supply chain disturbances and in the precision of the posterior estimates,

¹Shipping contracts typically span over a year, with route and schedule adjustments occurring every three to six months (Stopford, 2008; Meng et al., 2014).

especially when compared with alternative indices of global supply chain conditions.

The causal effects of supply chain disruptions. Using our identification strategy, we estimate a Bayesian SVAR with the ACR index under zero restrictions, yielding two key findings.

First, a supply chain shock leads to a surge in spare capacity, proxied by the import-weighted average spare capacity rate of the top five U.S. trading partners (Mexico, Canada, China, Germany, and Japan), which together account for more than half of U.S. goods imports. At the same time, product market tightness falls sharply, proxied by the imbalance between U.S. manufacturers' new orders and the U.S. dollar value of the import-weighted average spare capacity of the same trading partners. This shock also triggers a pronounced decline in U.S. real PCE and a persistent increase in the PCE price index, consistent with recent empirical evidence ([Khil and Weber, 2022](#); [Alessandria et al., 2023](#)). As predicted, productive capacity and supply chain shocks differ sharply in their effects on spare capacity and product market tightness: spare capacity (tightness) falls (rises) after a capacity shock but rises (falls) after a supply chain shock.

Second, the historical decomposition reveals four phases of U.S. headline PCE inflation since 2020. In the first phase (2020), the sharp disinflation was driven primarily by a contraction in aggregate demand and by a loosening of productive capacity constraints at the onset of the pandemic. In the second phase (2021 to mid-2022), inflation surged largely due to global supply chain disruptions. In the third phase (late 2022 to 2023), inflation eased as demand weakened, capacity improved, and supply chains recovered. In the final phase (2024 onward), inflation reaccelerated mildly, primarily reflecting supply chain factors, with demand and capacity continuing to exert a modest drag. A comprehensive set of sensitivity checks confirms that our results are robust across multiple dimensions, including the choice of proxy for global supply chain conditions, identification restrictions, model specification, and estimation method.

Related literature. Our study is related to several realms of research. As mentioned above, our model builds on [Barro and Grossman \(1971\)](#), [Michaillat and Saez \(2015, 2022\)](#), and [Ghassibe and Zanetti \(2022\)](#). It is also related to studies that focus on the effects of supply chain disturbances on output and inflation, using measures such as spare labor capacity ([Benigno and Eggertsson, 2023](#)), goods shortages ([Bernanke and Blanchard, 2025](#)), a quasi-kinked demand curve for produced goods ([Harding et al., 2023](#)), and capacity constraints ([Comín et al., 2023](#); [Merendino and Monacelli, 2025](#)).

Furthermore, our paper is related to work showing that transportation costs are important

for international trade and economic activity (Allen and Arkolakis, 2014; Brancaccio et al., 2020; Dunn and Leibovici, 2023), infrastructure investment (Fuchs and Wong, 2022; Brancaccio et al., 2024), asset prices (Smirnyagin and Tsyvinski, 2022), working capital (Antràs, 2023; Kim and Shin, 2023), inflation expectations (Acharya et al., 2023; Binetti et al., 2024), the design of new taxes and pricing rules to offset distortionary effects on the transportation network (Brancaccio et al., 2023), the interlinks between oil shocks and trading externalities in the supply chain (Bai and Li, 2022; Li et al., 2022), and the economic effects of supply chain disruptions during the COVID-19 pandemic (Finck and Tillmann, 2022; Ascoli et al., 2024; Finck et al., 2024).

The remainder of the paper is organized as follows. Section 2 constructs the ACR index measuring the state of the global supply chain. Section 3 develops our theoretical model and the identification restrictions for structural shocks. Section 4 presents the estimation results from a Bayesian SVAR. Section 5 concludes. An appendix provides further details, and our data are available on our website: <https://globalportcongestion.github.io/blog/intro.html>.

2. Measuring the State of the Global Supply Chain

This section constructs an index of global supply chain conditions by measuring congestion at major ports using satellite data on container-ship positions, speeds, and headings. We begin by explaining why we focus on containerized trade and by outlining key features of the industry. We then describe our satellite data, motivate the use of port congestion as a measure of global supply chain conditions, and present the algorithm that implements our approach. We conclude by reporting our results and comparing them with alternative indices in the literature.

2.1. Containerized Seaborne Trade: Some Basic Facts

Containerized seaborne trade plays a central role in the global supply chain, accounting for about 46% of all international trade.² In the U.S., container shipping carries more tonnage (nearly one billion short tons) and value (more than 0.7 trillion dollars) than any other transport mode, representing over 50% of U.S. trade by weight and roughly 30% by value (Bureau of Transportation Statistics, 2021). Although some high-value items, such as computer chips, are shipped by air, these products depend on other components, such as motherboards or hard drives, that move by container ship.

²See Notteboom et al. (2022). Most of the remainder consists of bulk cargo (e.g., oil, grain, ore, and coal) or specialized vessels (e.g., roll-on/roll-off ships for wheeled cargo).

As Brancaccio et al. (2020, p.2) explain, “The transportation sector … can be split into two categories: those that operate on fixed itineraries, much like buses, and those that operate on flexible routes, much like taxis. Container ships … belong to the first group.” These fixed itineraries center on seaports that act as hubs for freight collection and distribution. Even mild congestion at these ports can disrupt the tight schedules of supply chains and trade flows, leading to significant delay costs.

Before 2020, waiting times at ports were typically only a few hours. Disruptions linked to the COVID-19 pandemic, however, produced long delays, with waiting times of 2–3 days at several major ports and heavy financial losses.³ Since nearly 80% of world trade is shipped indirectly and the average shipment stops at five intermediate ports before reaching its final destination (Ganapati et al., 2024), the interconnected nature of global trade greatly amplifies delays from port congestion.⁴

The industry is also surprisingly concentrated. In 2022, there were only 5,589 container ships worldwide, of which roughly 500 belonged to the larger size classes.⁵ Hence, delays affecting even a single large ship can have significant consequences for global trade. For instance, the MSC Loreto, a new ultra-large container vessel, carries about 24,346 TEUs (twenty-foot equivalent units), each with a maximum cargo of 21,600 kilograms. At full capacity, it can load up to 240 thousand tons of cargo. A historical comparison highlights the scale: perhaps the most famous convoy of the Battle of the Atlantic during World War II, ONS 5, sailed from Liverpool to Halifax from April 29 to May 6, 1943, and became the center of an epic battle against 43 German U-boats. ONS 5 involved 49 merchant ships with a combined cargo capacity of roughly 219 thousand tons, about 10% less than that of the Loreto. Any delay in the Loreto’s loading or unloading has ramifications for tens of thousands of firms worldwide.

The escalation of port congestion during the COVID-19 pandemic was driven by multiple shocks. Delays arose from both upstream and downstream disruptions, including mobility re-

³Buyers and sellers of goods faced lower transport efficiency, higher operating costs, demurrage and detention charges, and difficulty meeting contractual obligations and market demand. For shippers and freight forwarders, delays were compounded by surcharges such as the port congestion surcharge (PCS), with fees climbing up to \$1,250 per container. Given that the average value of goods in a 40-foot container —the most common type— in 2020 was about \$109,000, the PCS alone represented a significant cost.

⁴For perspective on shipments to the U.S., it takes about 28 days to move a container from Shanghai to Los Angeles along the Trans-Pacific route (Freightos, 2024). A typical shipment stops at the Port of Ningbo-Zhoushan in China, Kaohsiung in Taiwan, Busan in South Korea, and Tokyo in Japan before arriving in Los Angeles. Assuming waiting times of 2–3 days at each intermediate port, the total delay amounts to 8–12 days, even before accounting for congestion at the origin and destination ports.

⁵See <https://unctad.org/rmt2022> (accessed December 29, 2023).

strictions from stay-at-home orders, port quarantine measures that reduced handling efficiency, late truck arrivals due to highway controls, and unopened containers at inland factories where workers were unavailable. A surge in demand for tradable goods in 2021 intensified congestion, as shipping companies deployed additional capacity on major routes (e.g., the Trans-Pacific corridor between East Asia and North America) and increased ship calls to meet demand (see Appendix A and [Bai et al., 2025a](#) for further details).

In short, the state of the global supply chain reflects the effects of multiple shocks. Our theoretical model, developed later in the paper, provides the identification assumptions needed to separate demand shocks from supply chain shocks that affect this system. Before turning to that model, however, we describe how our index of port congestion is constructed.

2.2. AIS Data

We use satellite data from the AIS, a real-time tracking system mandated by the IMO. International vessels exceeding 300 gross tons must carry an AIS transceiver that broadcasts ship information ([Heiland et al., 2025](#)). Each record includes the IMO number, timestamp, draft, speed, heading, and geographical coordinates.⁶ The AIS processes over 2,000 reports per minute and updates as often as every two seconds, providing comprehensive coverage of vessel movements worldwide from January 2016 to March 2025.⁷ Data on ship positions, speeds, and headings allow us to track movements within and across port zones.

2.3. A Density-Based Spatial Clustering Algorithm

The literature on maritime economics identifies port congestion as a key indicator of global supply chain conditions ([Cerdeiro and Komaromi, 2020](#); [Karimi-Mamaghan et al., 2020](#); [Bai et al., 2023](#); [Brancaccio et al., 2024](#)). A common approach measures congestion by estimating the probability that a vessel first moors in an anchorage area before docking at a berth ([Talley, 2009](#); [Talley and Ng, 2016](#); [Komaromi et al., 2022](#)). An anchorage is a location within a port where ships lower their anchors, while a berth is a designated site where vessels moor to load and unload cargo. In the absence of congestion, ships would proceed directly to a berth upon arrival.

⁶The draft measures the vertical distance from the bottom of a vessel’s keel to the water surface, indicating how deeply the ship is submerged.

⁷More than 99% of international container shipments are carried by vessels above 500 gross tons. Even smaller vessels (below 300 gross tons) usually carry AIS transceivers because of their substantial safety benefits at a relatively low cost (about \$1,000 for a basic unit). Our coverage is therefore nearly universal.

Measuring congestion, therefore, requires identifying berth and anchorage areas, a task that previous studies have typically undertaken using nautical charts of individual ports. This approach is labor-intensive and difficult to generalize across heterogeneous global port layouts.⁸ To overcome these limitations, we develop a spatial clustering algorithm that accurately distinguishes areas within ports and applies to ports with diverse morphologies.

Our algorithm identifies distinct port areas by analyzing the density of container ships' mooring points in the AIS data through two layers of clustering. The first layer detects high-density regions (locations with many AIS observations) and treats them as potential berths or anchorages. The second layer distinguishes between them using domain knowledge of vessel behavior in port. When vessel headings are orderly and closely aligned, the area is classified as a berth; when headings are more dispersed or irregular, it is classified as an anchorage.

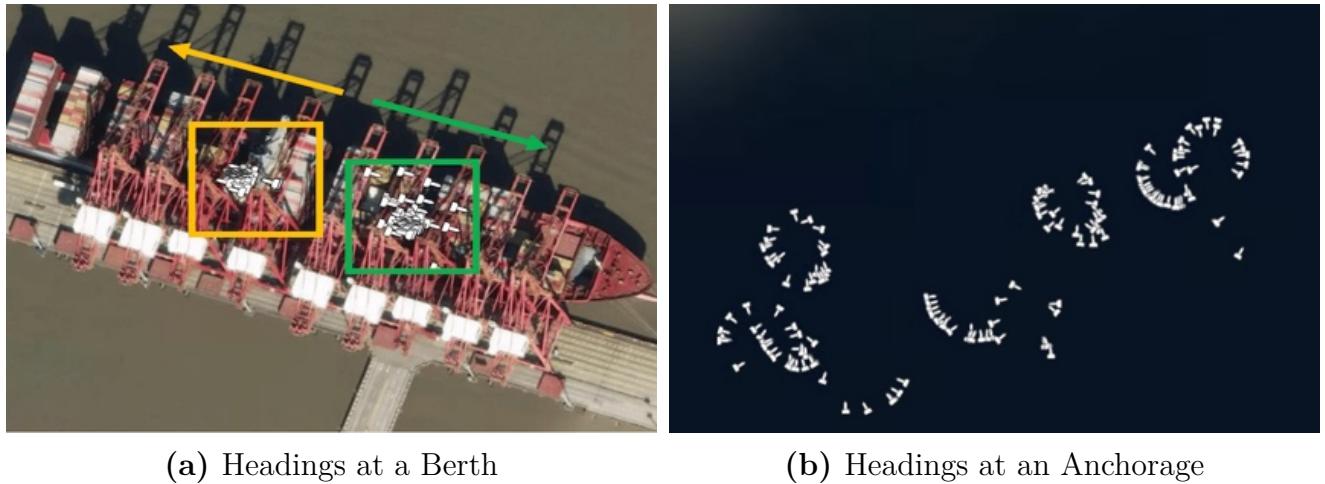


Figure 1: Information on Headings: Two Examples

Figure 1 illustrates this distinction. The left panel shows two clusters, one in orange and one in green. Both contain many AIS observations, with bows (the thin tips of the white signs) closely aligned —pointing left in the orange cluster and right in the green cluster— representing two different mooring headings. Superimposing these clusters on a satellite image confirms that they correspond to a berth. The right panel, by contrast, shows several clusters where vessel headings

⁸Nautical charts have three main drawbacks: (i) they are static and rarely updated, so new berths may be missing; (ii) the vast number of ports makes manual delineation of boundaries nearly impossible; and (iii) they do not capture with precision the heterogeneity of areas within ports. Thus, studies relying on nautical charts usually focus on one or a few ports (Chen et al., 2016; Feng et al., 2020). Inspecting satellite images can help identify berths with fixed locations and boundaries —for example, Appendix C.2 validates our estimates for the geographically complex Port of Ningbo-Zhoushan, south of Shanghai. However, satellite images generally cannot identify anchorages, whose locations and boundaries vary with weather and port conditions.

are random, with some forming a ring shape, consistent with anchorage behavior. Appendix C provides further details on the algorithm, including pseudocode and a case study of the Port of Ningbo-Zhoushan.

Our algorithm addresses two key challenges faced by existing clustering methods. First, it adapts to variability in ship mooring densities across ports —arising from differences in trade volume, vessel frequency, and geographical features— by iteratively refining clustering parameters for each port. Second, through its two-layer design, the algorithm distinguishes between berths and anchorage areas even in dense mooring environments. It is also adaptable to other applications, such as analyzing port efficiency, canal traffic, and stress at maritime choke points. More broadly, its core mechanism, transforming domain knowledge into non-spatial attributes for iterative clustering, provides a versatile tool for classifying clusters with specific labels in diverse contexts (e.g., disease hotspots, urban planning, environmental monitoring).

Figure 2 shows that our algorithm accurately identifies anchorage and berth areas across ports with diverse geographical and operational conditions. In each panel, we superimpose the identified anchorage areas (colored red, yellow, blue, purple, pink, cyan, and orange) and berth areas (markers of other colors) on satellite images of four major container ports: Ningbo-Zhoushan (Panel a), Los Angeles and Long Beach (Panel b), Rotterdam (Panel c), and Singapore (Panel d). Separate figures for anchorages and berths in each port are provided in Appendix C.2.

2.4. The ACR Index

Port congestion arises when ships cannot immediately load or unload cargo upon arrival at ports and must wait in an anchorage area until a berth becomes available. For the top 50 container ports worldwide, denoted as \mathcal{P} , we count the number of delayed ship visits to each port p where a ship first moors in an anchorage before docking at a berth.⁹ These top 50 ports represent 75.6% of total containerized trade worldwide.¹⁰ We then calculate the congestion rate for each port p by dividing the number of delayed ship visits by the total number of ship visits:

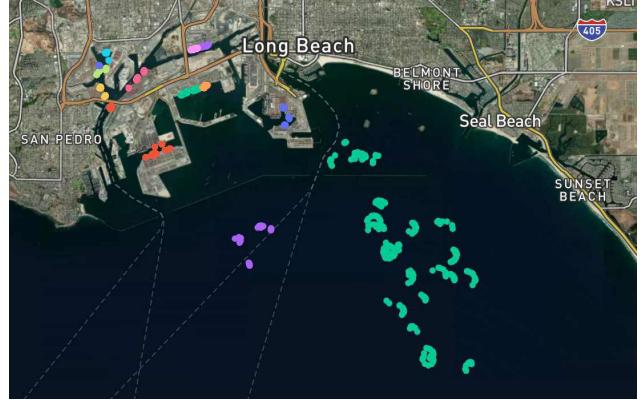
$$\text{Congestion}_{p,t} \equiv \frac{\text{Delayed}_{p,t}}{\text{Delayed}_{p,t} + \text{Undelayed}_{p,t}}, \quad \forall p \in \mathcal{P}, \quad (1)$$

⁹A ship visit is classified as delayed if at least one AIS observation shows zero speed and coordinates within the anchorage area —as mapped by the geographical boundaries identified by our clustering algorithm— before the ship docks at a berth.

¹⁰Extending the index to more ports is straightforward but computationally costly.



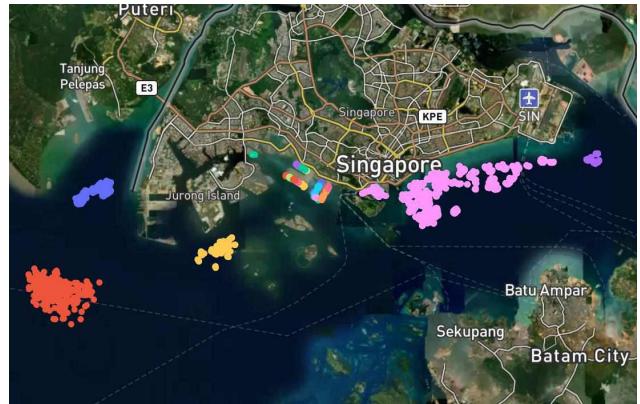
(a) Ningbo-Zhoushan, China



(b) Los Angeles and Long Beach, U.S.



(c) Rotterdam, Netherlands



(d) Singapore

Figure 2: Identification of Anchorage and Berth Areas of a Port Using Machine Learning

Note. Each panel is based on the first 50,000 AIS observations of container ships entering each of the four major ports since January 1, 2020.

where $\text{Delayed}_{p,t}$ and $\text{Undelayed}_{p,t}$ denote the number of delayed and undelayed ship visits at port p in month t , respectively.¹¹

We choose a rate measure of port congestion over the commonly used time measure (Fuchs and Wong, 2022; Brancaccio et al., 2024) because it is unaffected by large differences in vessel wait times across ports. For example, in June 2022, the average anchorage wait time in Savannah, U.S., was 186 hours, while in Shanghai, China, it was virtually zero. Moreover, the rate measure

¹¹A vessel's draft generally reflects its cargo load, suggesting that we could weight delays by the volume of cargo affected. However, unlike bulk carriers or oil tankers —whose drafts fluctuate sharply between voyages as they load only after discharging previous cargoes— container ships load and unload simultaneously, generating minimal draft variation. As a result, weight-augmented indices would closely mirror our original congestion indices.

enables more reliable cross-port comparisons, as local anchorage times may reflect port size and infrastructure conditions rather than supply chain disruptions.¹²

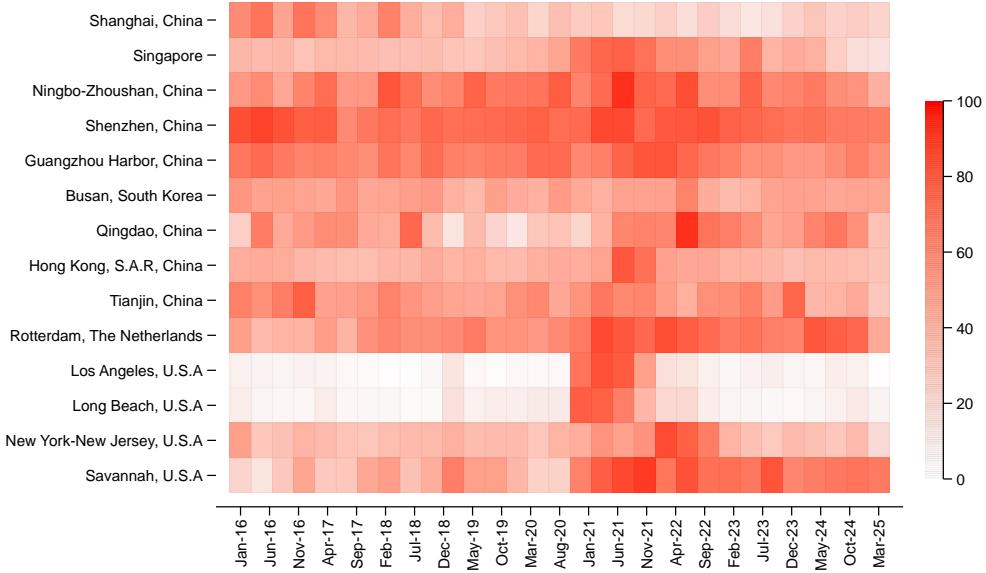


Figure 3: Congestion Rates for the Major Container Ports Worldwide

Notes. Heatmap of monthly congestion rates for the top ten global container ports plus the Ports of Los Angeles, Long Beach, New York–New Jersey, and Savannah, from January 2016 to March 2025. The congestion rate for each port is normalized and expressed as a percentage of its peak value within the sample period. Darker shades indicate higher congestion levels, as defined in Equation (1), for the respective port in a given month.

Figure 3 shows monthly congestion rates for the top ten ports worldwide, along with the Ports of Los Angeles, Long Beach, New York–New Jersey, and Savannah (which together handle more than 60% of U.S. containerized imports), from January 2016 to March 2025. While we rely on data from the top 50 ports worldwide in the analysis below, the ports in Figure 3 account for more than 30% of global containerized seaborne trade and, thus, summarize our main findings.

Our data indicate that the onset of the COVID-19 pandemic in March 2020 had limited immediate effects. The situation changed in the fall of 2020, when congestion rates rose across most major ports, as reflected in the deepening shades in Figure 3 after August 2020. The intensification continued through 2021 and early 2022, with some ports experiencing historically high congestion levels. For example, nearly 80% of inbound ships at the Port of Los Angeles were

¹²Nevertheless, Appendix H.3 also constructs an index of Average Congestion Time (ACT). Consistent with Brancaccio et al. (2024), the ACT index measures the average number of hours a container ship waits in a port’s anchorage before docking at a berth, weighted by the relative number of ship visits to each port. Using the ACT index in the causality assessment yields results that are quantitatively similar to those obtained with our original ACR index.

unable to dock immediately upon arrival in late 2020.¹³

From late 2022 onward, congestion pressures gradually eased across most ports. By 2023, chromatic intensity in the heatmap diminished sharply, indicating a normalization of shipping. Some residual congestion persisted, however, in ports such as Shenzhen, Rotterdam, and Savannah, where elevated rates continued into 2023. By early 2025, congestion levels had stabilized at much lower levels than during the pandemic peak, though not uniformly across all ports.

Subsequently, we define the Average Congestion Rate (ACR) as the weighted average of congestion rates across the top 50 container ports, where the weights correspond to the relative number of ship visits to each port:

$$\text{ACR}_t \equiv \sum_{p \in \mathcal{P}} \left[\frac{\text{Delayed}_{p,t} + \text{Undelayed}_{p,t}}{\sum_{p \in \mathcal{P}} (\text{Delayed}_{p,t} + \text{Undelayed}_{p,t})} \cdot \text{Congestion}_{p,t} \right],$$

with $\text{Congestion}_{p,t}$ defined in Equation (1).¹⁴



Figure 4: ACR Index

Note. The ACR index is computed as the weighted average of congestion rates across the top 50 container ports worldwide (as of June 15, 2022), using the relative number of ship visits as weights. The index is expressed in percentage terms and has been seasonally adjusted.

Figure 4 plots the ACR index. Before 2018, the index declined steadily and then stabilized

¹³This estimate aligns with official statistics. According to the Pacific Merchant Shipping Association, the share of container ships in Los Angeles waiting five or more days for unloading rose from 10% in August to 26% in December 2020.

¹⁴We use the relative number of ship visits to each port as weights in the ACR index, as they capture each port's importance in the global supply chain. For example, a slight increase in the congestion rate at the Port of Hong Kong may signal a larger disruption than a large increase at the Port of Manila.

around 28%, reaching a minimum of 25% from early 2019 to mid-2020. This pattern reflects the substantial investments made worldwide in earlier years to expand port capacity and strengthen supply chain resilience. Conditions shifted sharply thereafter: the index climbed throughout 2020 and 2021, peaking at 37% in June 2021, a clear sign of the severe strains the COVID-19 pandemic imposed on global supply chains. At that point, nearly one in every three container ships entering one of the top 50 ports encountered delays due to congestion. Although the index stayed elevated until mid-2022, it began to fall later that year and returned to the sample median (30%) by mid-2023. By then, port congestion had largely normalized, and global supply chains were operating more smoothly, even if levels remained slightly above the pre-COVID average.

This period of relative stabilization, however, proved short-lived. Geopolitical disruptions to global shipping routes in 2024 pushed the index back above 32%. Attacks on vessels in the Red Sea and drought-related restrictions on Panama Canal transits forced widespread rerouting, caused schedule disruptions, and led to bunching at alternative ports (Bai et al., 2025a). By early 2025, however, the index dropped sharply below 25%. This decline reflects the gradual normalization of shipping routes, the easing of Panama Canal restrictions, and softer global trade volumes following newly imposed U.S. tariffs that reduced import demand. In sum, the ACR index traces the sharp swings in global port congestion—from its COVID-19 peak, to post-pandemic normalization, to renewed geopolitical pressures in 2024, and finally the marked easing in early 2025.

2.5. Discussion

Several aspects of our index merit further discussion, as they underpin the modeling assumptions in Section 3, the identification restrictions in Section 4, and the sharper identification they yield in causality analyses of supply chain disruptions relative to alternative indices.

Short-run rigidities in containerized shipping. Containerized shipping entails two short-run frictions: an *economic* margin and an *operational* margin. On the economic side, service contracts fix invoiced freight rates for at least a one-month horizon, so the maximum transportation cost a match can bear without separation—the “reservation” transportation cost—remains rigid in the short run when profitability does not change. On the operational side, fixed rotations, berth windows, and alliance commitments make schedules and arrivals similarly rigid, as decisions by shipping companies are typically revised only every three to six months (Stopford,

2008; Meng et al., 2014).¹⁵

Supply chain disturbances (e.g., mobility restrictions, quarantines, and related disruptions discussed in Subsection 2.1) reduce effective service rates at ports on impact, lengthening queues and raising delays while leaving both contracts and schedules unchanged within a month. These delays increase transportation costs through congestion surcharges and premiums on expedited services, while the reservation transportation cost remains rigid because trade volumes and profitability do not adjust immediately.¹⁶ In the model developed in the next section, we represent a supply chain disturbance as an unexpected, broad increase in transportation costs, with the reservation threshold remaining rigid in the short run.

By contrast, demand or capacity shocks operate primarily through quantities and profitability within a month. They alter firms' willingness to pay and thus shift the reservation transportation cost, even though invoiced rates remain on contract. However, because schedules and arrivals are still fixed over this horizon, these shocks do not change congestion on impact. This distinction yields two short-run implications that we exploit empirically and theoretically: (i) under supply chain disturbances, the ACR rises and realized schedules lengthen (through queues and bunching) while the contractual price cap continues to bind the reservation threshold, whereas (ii) under demand or capacity shocks, the ACR and realized schedules show no immediate response even as willingness to pay and the reservation threshold adjust.

In Section 4, we leverage the distinct timing of the ACR response, together with restrictions on macroeconomic aggregates, to isolate operational supply chain disturbances from demand or capacity shocks. We validate the short-run operational rigidity of containerized shipping by showing (i) no statistically significant on-impact response of the ACR index to a Bauer–Swanson monetary policy shock (Bauer and Swanson, 2023) —reflecting demand changes— using local projections (LPs) in Appendix B, and (ii) no statistical correlation between port congestion and oil price movements —reflecting capacity changes— in Appendix A, despite the fact that higher oil prices directly incentivize slower steaming.

Alternative indices. Subsection 4.4 compares our ACR index with other popular measures

¹⁵The same operational rigidity is evident in the industry's "hurry up and wait" practice: even with forewarnings of downstream delays, vessels rarely alter routes or speeds on short notice because berth windows, feeder connections, alliance slot commitments, and contractual terms jointly constrain flexibility. See Appendix A for details.

¹⁶Shipping prices are also highly sensitive to physical congestion due to the small number of container ships: even a few delays drive up freight rates.

of global supply chain conditions, including the Harper Peterson Time Charter Rates Index (HARPEX), the New York Fed’s Global Supply Chain Pressure Index (GSCPI), and the Supply Disruptions Index (SDI) from [Smirnyagin and Tsyvinski \(2022\)](#) and [Liu et al. \(2024\)](#).¹⁷ Relative to these alternatives, supply chain shocks identified with the ACR produce impulse responses of the U.S. PCE price index that are larger in magnitude, more precisely estimated, and in line with theoretical predictions. Moreover, variance decompositions indicate that when the ACR is used, supply chain shocks explain a larger share of the forecast error variance of prices at medium horizons, whereas demand shocks dominate under alternative proxies. Together, this evidence indicates that the ACR index provides a more *accurate* proxy for global supply chain conditions.

Appendix H.4 also develops a targeted ACR index for major ports along the Trans-Pacific route, the key shipping lane between East Asia (mainly China) and the U.S. Using this index yields causality results nearly identical to those obtained with the global ACR index.

Finally, integrating high-frequency AIS data with our spatial clustering algorithm allows us to construct congestion indices at frequencies higher than the monthly level. Appendix C.3 reconstructs the ACR and ACT indices using weekly AIS data and shows that although volatility increases, the weekly series exhibit patterns consistent with their monthly counterparts.

3. A Model of the Global Supply Chain

Next, we develop a theoretical model of the global supply chain to establish the identification restrictions for our causality analysis. Our economy consists of producers, retailers, and households. Producers manufacture goods using a fixed-factor endowment and incur transportation costs when selling goods to retailers. Retailers purchase goods from producers but face search frictions that make it difficult to meet with them. Retailers then sell goods to households. Households own both producers and retailers, accruing all profits generated through these trades.

We separate producers and retailers by location to capture the need for firms to trade within a global supply chain, where search frictions make such trade non-trivial. In addition, transportation costs and search frictions limit the allocative role of prices in clearing the quantity of goods sold from producers to retailers. Finally, the model is not designed to study port congestion *per se*, but rather supply chain disturbances, of which rising port congestion is a manifestation.¹⁸

¹⁷ Appendix H details this comparison and shows that differences across indices significantly affect how supply chain disruptions and their macroeconomic consequences are interpreted.

¹⁸ Appendix D discusses the evidence of search frictions in the goods market and the relevance of transportation

Our model allows us to analyze three shocks that influence the ACR index: an aggregate demand shock (e.g., changes in household money holdings driven by monetary policy or shifts in preferences for consumption), a productive capacity shock (e.g., changes in producers' fixed-factor endowment), and a supply chain shock (e.g., a broad increase in transportation costs). It delivers distinct predictions for how spare productive capacity and product market tightness co-move with prices and consumption in response to each shock, thereby enabling the unique identification of the causal effects of supply chain disruptions.

3.1. Producers and Retailers

There is an exogenous unit mass of producers and an endogenous measure of retailers. When matched with a retailer, a producer manufactures $y = l$ final goods using its fixed-factor endowment.¹⁹ Producers sell goods to retailers in a frictional product market that prevents the sale of full capacity. Each unmatched retailer (identified by the subscript U) makes one visit per period to an unmatched producer, with each visit entailing a fixed cost per unit of the final good $\rho > 0$. When a producer and a retailer meet and trade (as discussed below), the retailer resells the purchased good to the household at a price p .

Matching process. In each period, the number of meetings (\mathcal{M}) between unmatched producers, x_U , and retailers, i_U , is governed by a constant-returns-to-scale matching function:

$$\mathcal{M} = (x_U^{-\xi} + i_U^{-\xi})^{-\frac{1}{\xi}},$$

where ξ is the elasticity of substitution between x_U and i_U . We assume $\xi > 0$ such that $\mathcal{M} \leq \min \{x_U, i_U\}$, which is a necessary property of a matching function.

Product market tightness $\theta \equiv i_U/x_U$ is the ratio between the number of visits by unmatched retailers and the number of unmatched producers. Individual firms take product market tightness as given. The probability that a producer meets a retailer is:

$$f(\theta) = \frac{\mathcal{M}}{x_U} = (1 + \theta^{-\xi})^{-\frac{1}{\xi}}, \quad (2)$$

costs for the severance of commercial trade.

¹⁹We abstract from modeling producers' endogenous production decisions, as this would require a multi-country, multi-sector production network model to study the transmission of sectoral supply disruptions across countries (di Giovanni et al., 2022, 2023).

and the probability that a retailer meets a producer is

$$q(\theta) = \frac{\mathcal{M}}{i_U} = (1 + \theta^\xi)^{-\frac{1}{\xi}}.$$

The function $f(\theta)$ satisfies $f(0) = 0$, $\lim_{\theta \rightarrow +\infty} f(\theta) = 1$, and $f_\theta(\theta) > 0$, while $q(\theta)$ satisfies $q(0) = 1$, $\lim_{\theta \rightarrow +\infty} q(\theta) = 0$, and $q_\theta(\theta) < 0$. Three additional properties that will be useful later are $f(\theta)/q(\theta) = \theta$, $f_\theta(\theta) = q(\theta)^{1+\xi}$, and $f(\theta)^\xi + q(\theta)^\xi = 1$.

Transportation cost. Producers pay a per-unit idiosyncratic transportation cost to ship their goods to retailers.²⁰ In each period, producers draw a per-unit transportation cost z from the log-normal distribution $G(z)$ with scale parameter γ and shape parameter σ , that is, $G(z) \equiv \Phi[(\ln z - \gamma)/\sigma]$, where $\Phi(\cdot)$ is the standard normal c.d.f.²¹ As discussed later, there exists a reservation transportation cost \bar{z} above which matches are unprofitable and severed ($z > \bar{z}$), while they continue otherwise ($z \leq \bar{z}$).

Value functions. At the beginning of each period, matched producers sell the manufactured goods to retailers and pay the transportation costs. Matched retailers sell their purchased goods to households and pay the wholesale price of goods to producers. Unmatched producers and retailers search to form a match with each other. At the beginning of the next period, each producer draws a new transportation cost, and the match continues if the new cost is sufficiently low to generate a positive surplus from trade.

Four value functions characterize the returns associated with the different statuses of producers and retailers. The value for a matched producer (denoted by the subscript M), $X_M(z)$, is

$$X_M(z) = (r(z) - z)l + \beta \mathbb{E}_{z'} [\max(X_M(z'), X_U)], \quad (3)$$

where $r(z)$ is the endogenous wholesale price per unit of the final good, β is the discount factor, and z' is the transportation cost drawn at the beginning of the next period. Equation (3)

²⁰Our results hold if the transportation cost is borne by retailers instead because the match separation condition (11) is invariant to this modeling choice. For simplicity, we also assume that the household receives this shipping cost as payment for its work in moving the goods.

²¹The assumption that transportation costs follow a log-normal distribution aligns with [Kasahara and Lapham \(2013\)](#). A random distribution captures heterogeneity of transportation costs and generates an extensive margin of trade: only matches with $z \leq \bar{z}$ are viable, while those with $z > \bar{z}$ are severed. This feature would be absent under a fixed cost, where all matches survive, and comparative statics would operate only through intensive adjustments (i.e., changes in quantities within existing matches rather than changes in the fraction of surviving matches). The stochastic specification also provides a tractable way to study shocks to the cost distribution (through the scale parameter γ). Alternatively, one could model a full transportation sector where interactions among producers, shipping companies, and retailers determine costs, as in [Brancaccio et al. \(2020\)](#), [Bai and Li \(2022\)](#), [Dunn and Leibovici \(2023\)](#), and [Bai et al. \(2025a\)](#), but we keep the current setting for tractability.

shows that the present value of being a matched producer consists of the current profit margin, $(r(z) - z) l$, plus a continuation value that depends on whether the producer separates from the match. Separation is determined by the next period's transportation cost z' , with the max operator selecting the optimal continuation or separation decision.

The value for an unmatched producer, X_U , is

$$X_U = \beta f(\theta) \mathbb{E}_{z'} [\max(X_M(z'), X_U)] + \beta (1 - f(\theta)) X_U. \quad (4)$$

With probability $f(\theta)$, an unmatched producer meets a retailer and then decides whether to separate if the new transportation cost makes the match unprofitable. With probability $1 - f(\theta)$, the producer fails to meet a retailer and remains unmatched at the start of the next period.

The value for a matched retailer, $I_M(z)$, is

$$I_M(z) = (p - r(z)) l + \beta \mathbb{E}_{z'} [\max(I_M(z'), I_U)]. \quad (5)$$

The retailer earns the resale price p from households for each unit of the good and pays the wholesale price $r(z)$ to the producer. As before, the max operator selects the optimal continuation or separation decision, conditional on z' .

If the realized transportation cost makes the match unprofitable, the retailer separates and begins the next period with value:

$$I_U = -\rho l + \beta q(\theta) \mathbb{E}_{z'} [\max(I_M(z'), I_U)] + \beta (1 - q(\theta)) I_U, \quad (6)$$

where ρ is the fixed cost per unit of the final good that the retailer incurs with each visit to a producer. Free entry into the product market drives the value for an unmatched retailer to zero; that is, $I_U = 0$.

Nash bargaining. The total surplus from matching is equal to:

$$S(z) = X_M(z) - X_U + I_M(z) - I_U, \quad (7)$$

and it is split through Nash bargaining. The producer receives a constant share η of the total surplus, while the retailer receives the remaining share $1 - \eta$, implying:

$$\eta (I_M(z) - I_U) = (1 - \eta) (X_M(z) - X_U). \quad (8)$$

Given the Nash bargaining rule (8), the value functions (3), (4), (5), and the free-entry condition $I_U = 0$, the wholesale price that splits the surplus is

$$r(z) = \eta(p + \rho\theta) + (1 - \eta)z. \quad (9)$$

The wholesale price is a weighted average of the retail price (adjusted for the search cost $\rho\theta$) and the transportation cost z , with the weights determined by bargaining power. When producer bargaining power is low ($\eta \rightarrow 0$), the wholesale price approaches the transportation cost z . In addition, higher tightness, which worsens retailers' bargaining positions by lowering their matching probability, increases the wholesale prices they pay to producers.

Match separation. Since the total value for a matched producer and a matched retailer, that is, $X_M(z) + I_M(z)$, strictly decreases with the transportation cost z , there exists a cut-off transportation cost \bar{z} above which matches become unprofitable and are therefore severed. This cut-off makes the total surplus in Equation (7) equal to zero:

$$S(\bar{z}) = 0. \quad (10)$$

Substituting the value functions (3), (4), (5) and the free-entry condition $I_U = 0$ into Equation (10) yields the match separation condition, expressed as a function of p , \bar{z} , and θ :

$$\mathbb{F}(p, \bar{z}, \theta) = (p - \bar{z})l + (1 - \eta f(\theta)) \beta \mathbb{E}_{z'} S(z') = 0, \quad (11)$$

where $\mathbb{E}_{z'} S(z') = \int_0^{\bar{z}} S(z') dG(z')$ is the expected surplus.

Match creation. Using the value function for an unmatched retailer in Equation (6), together with the free-entry condition $I_U = 0$, the match creation condition can be written as a function of \bar{z} and θ :

$$\mathbb{H}(\bar{z}, \theta) = \frac{\rho l}{q(\theta)} - (1 - \eta) \beta \mathbb{E}_{z'} S(z') = 0. \quad (12)$$

Aggregate supply. Aggregate supply in the economy is the quantity of goods traded by retailers and producers that survive separation for a given productive capacity, equal to the producers' total factor endowment l . To determine this supply, we consider the law of motion for the number of matched producers at the beginning of the next period, x'_M :

$$x'_M = G(\bar{z})x_M + f(\theta)G(\bar{z})x_U,$$

and for the number of unmatched producers:

$$x'_U = [1 - f(\theta) + f(\theta)(1 - G(\bar{z}))] x_U + (1 - G(\bar{z})) x_M.$$

Using the identity $x_M + x_U = 1$, the law of motion for matched producers can be rewritten as:

$$x'_M = f(\theta)G(\bar{z}) + (G(\bar{z}) - f(\theta)G(\bar{z})) x_M. \quad (13)$$

Aggregate supply is, thus, the quantity of goods provided by matched producers for a given productive capacity:

$$c_s(\bar{z}, \theta) = x_M(\bar{z}, \theta)l. \quad (14)$$

3.2. The Representative Household

The representative household derives utility from consuming goods and holding real money balances:

$$u\left(c, \frac{m}{p}\right) = \frac{\chi}{1+\chi} c^{\frac{\varepsilon-1}{\varepsilon}} + \frac{1}{1+\chi} \left(\frac{m}{p}\right)^{\frac{\varepsilon-1}{\varepsilon}},$$

where c denotes consumption, m nominal money balances, p the price level, $\chi > 0$ the relative preference for consumption over money, and $\varepsilon > 1$ the elasticity of substitution between consumption and real money balances. Following [Michaillat and Saez \(2015\)](#), we adopt this specification to ensure that aggregate demand drives fluctuations in macroeconomic aggregates.

The household owns both producers and retailers, receiving lump-sum rebates from their profits as well as compensation for transportation costs incurred in moving goods. Taking prices as given, the household chooses consumption and nominal money balances to maximize utility, subject to the budget constraint:

$$pc + m \leq \mu + \underbrace{pc_s(\bar{z}, \theta) - \rho l i_U}_{\text{Profits of Producers \& Retailers}} - \underbrace{\int_0^{\bar{z}} z' c_s(\bar{z}, \theta) dG(z')}_{\text{Transportation Costs}} + \underbrace{\int_0^{\bar{z}} z' c_s(\bar{z}, \theta) dG(z')}_{\text{Transportation Costs}},$$

where $\mu > 0$ is the household's endowment of nominal money, $c_s(\bar{z}, \theta)$ is the aggregate supply as defined in Equation (14), and $\rho l i_U$ represents the aggregate search outlay borne by unmatched retailers. Solving the household's optimization problem yields the optimality condition:

$$\frac{\chi}{1+\chi} c^{-\frac{1}{\varepsilon}} = \frac{1}{1+\chi} \left(\frac{m}{p}\right)^{-\frac{1}{\varepsilon}}. \quad (15)$$

Aggregate demand. Aggregate demand is the level of consumption that maximizes utility at a given price when the money market clears, a condition that holds both in and out of equilibrium. Substituting m with μ in Equation (15) and rearranging yields:

$$c_d(p) = \chi^\varepsilon \frac{\mu}{p}, \quad (16)$$

which is strictly decreasing and convex in p on $(0, +\infty)$. Because a higher price reduces real money balances, aggregate demand declines with the price.

3.3. Equilibrium and Steady State

We are now ready to define a period equilibrium (i.e., an equilibrium for a given period) and a steady state.²²

Definition 1. *A (period) equilibrium for this economy is a price p , a reservation transportation cost \bar{z} , and a product market tightness θ such that the match separation condition (11) and the match creation condition (12) hold simultaneously:*

$$\mathbb{F}(p, \bar{z}, \theta) = \mathbb{H}(\bar{z}, \theta) = 0,$$

and the retailer–household market clears:

$$c_s(\bar{z}, \theta) = c_d(p),$$

where aggregate supply $c_s(\bar{z}, \theta)$ evolves according to the law of motion for matched producers in Equation (13).

Rather than analyzing the full transition dynamics of the equilibrium after a shock (this analysis is relegated to Appendix E.5), we focus on the steady state. As shown in Section 3.4, comparative statics suffice to derive the identification restrictions for each shock of interest in our causality analysis.

Setting $x'_M = x_M$ in Equation (13) gives the steady-state number of matched producers:

$$x_M^{ss}(\bar{z}, \theta) = \frac{f(\theta)G(\bar{z})}{1 - G(\bar{z}) + f(\theta)G(\bar{z})}.$$

Steady-state aggregate supply is the quantity of goods provided by the steady-state number of

²²We define a period equilibrium rather than a full sequential equilibrium to simplify notation. It can be viewed as the prices and allocation at time t given the history of past shocks.

matched producers, given productive capacity l :

$$c_s^{ss}(\bar{z}, \theta) = x_M^{ss}(\bar{z}, \theta)l = \frac{f(\theta)G(\bar{z})}{1 - G(\bar{z}) + f(\theta)G(\bar{z})} l. \quad (17)$$

The steady state is defined in Definition 2, and its unique existence is established in Proposition 1.

Definition 2. *The steady state of this economy is a price p^* , a reservation transportation cost \bar{z}^* , and a product market tightness θ^* that jointly satisfy the match separation condition $\mathbb{F}(p, \bar{z}, \theta) = 0$, the match creation condition $\mathbb{H}(\bar{z}, \theta) = 0$, and the retailer–household market clearing condition $c_s^{ss}(\bar{z}, \theta) = c_d(p)$, that is:*

$$\frac{f(\theta)G(\bar{z})}{1 - G(\bar{z}) + f(\theta)G(\bar{z})} l = \chi^\varepsilon \frac{\mu}{p}. \quad (18)$$

Proposition 1. *A unique steady state $(p^*, \bar{z}^*, \theta^*)$ exists in which the match separation condition, the match creation condition, and the retailer–household market clearing condition all hold simultaneously.*

Proof. See Appendix E.1. ■

Having established the existence and uniqueness of the steady state, we now define the steady-state aggregate supply schedule $p \mapsto c_s^{ss}(p)$ implied by the partial equilibrium in the producer–retailer market—i.e., when the match separation and creation conditions jointly hold, $\mathbb{F}(\bar{z}, \theta; p) = \mathbb{H}(\bar{z}, \theta) = 0$ for a given p —and examine its properties in the following proposition.

Proposition 2. *For any $\bar{z} \geq \bar{z}_{\min}$, where \bar{z}_{\min} satisfies:*

$$\int_0^{\bar{z}_{\min}} G(z') dz' = \frac{\rho}{(1 - \eta)\beta}, \quad (19)$$

the steady-state aggregate supply schedule $p \mapsto c_s^{ss}(p)$, arising from the partial equilibrium in the producer–retailer market, has the following properties:

1. *The mapping*

$$\bar{z} \mapsto p(\bar{z}) \equiv \bar{z} - (1 - \eta f(\bar{z})) \beta \int_0^{\bar{z}} G(z') dz',$$

is continuously differentiable and strictly increasing, where:

$$f(\bar{z}) = (1 - q(\bar{z})^\xi)^{\frac{1}{\xi}}, \quad q(\bar{z}) = \frac{\rho}{(1 - \eta)\beta \int_0^{\bar{z}} G(z') dz'}.$$

Consequently, there exists a unique, continuously differentiable steady-state aggregate supply schedule $p \mapsto c_s^{ss}(p)$, represented by the parametric curve $(p(\bar{z}), c_s^{ss}(\bar{z}))$ on $\bar{z} \in [\bar{z}_{\min}, +\infty)$;

2. $\lim_{p \rightarrow p_{\min}} c_s^{ss}(p) = 0$ and $\lim_{p \rightarrow +\infty} c_s^{ss}(p) = l$, where p_{\min} is defined as:

$$p_{\min} \equiv \bar{z}_{\min} - \beta \int_0^{\bar{z}_{\min}} G(z') dz' = \bar{z}_{\min} - \frac{\rho}{1-\eta}; \quad (20)$$

3. $c_s^{ss}(p)$ is strictly increasing in p on $[p_{\min}, +\infty)$, and converges to a constant as $p \rightarrow +\infty$;
4. $c_s^{ss}(p)$ is locally convex near p_{\min} if $\xi \in (0, 1)$, linear if $\xi \geq 1$, and strictly concave for sufficiently large p .

Proof. See Appendix E.2. ■

Proposition 2 shows that aggregate supply rises with price through two reinforcing mechanisms. First, a higher price raises the surplus from producer–retailer matches, strengthening retailers’ incentives to search and thereby increasing product market tightness and the probability of successful matches. As more matches are formed, aggregate supply expands. Second, a higher price raises the reservation transportation cost, allowing matches that would otherwise have been abandoned to continue, further boosting supply. Although search frictions and transportation costs constrain the flow of goods and generate spare capacity, the model preserves the standard positive relationship between price and aggregate supply.

Panel 5a in Figure 5 plots aggregate demand, aggregate supply, and the steady state in the downstream retailer–household market, identified by their intersection in the (c, p) -plane. The steady-state aggregate supply schedule c_s^{ss} (blue line) is upward-sloping, as in standard models. Spare capacity —measured as the difference between productive capacity and actual output, $l - c_s^{ss}$ — captures the combined effects of search frictions and transportation costs.

Panel 5b in Figure 5 complements this by depicting the match separation condition (11), the match creation condition (12), and the steady state in the upstream producer–retailer market, identified by their intersection in the (\bar{z}, θ) -plane.²³

²³For reference, Appendix E.3 derives the slope and curvature of the match separation and creation schedules. Under mild conditions, the match separation condition $\mathbb{F}(\bar{z}, \theta; p) = 0$ generates a strictly decreasing and strictly convex relationship between θ and \bar{z} , whereas the match creation condition $\mathbb{H}(\bar{z}, \theta) = 0$ yields a strictly increasing and strictly convex relationship.

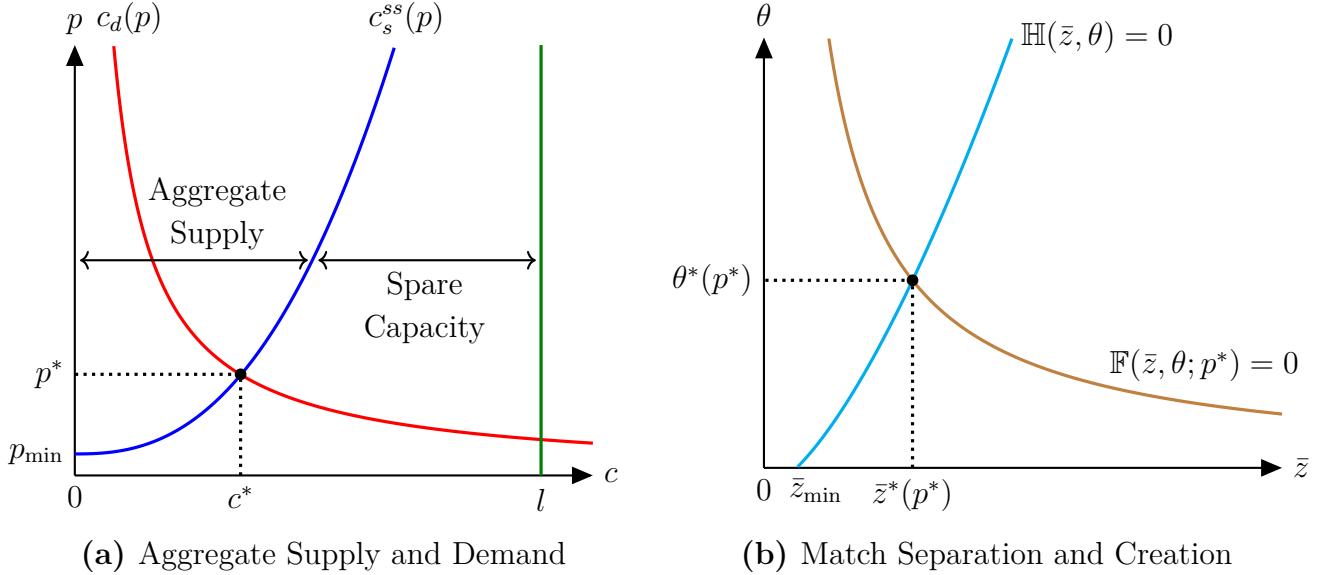


Figure 5: Graphical Representation of the Steady State

Notes. In Panel 5a, $c_d(p)$, $c_s^{ss}(p)$, and l represent the aggregate demand schedule, the steady-state aggregate supply schedule, and productive capacity, respectively. The threshold p_{\min} is defined in Equation (20). Panel 5b shows the partial equilibrium in the producer–retailer market at the steady-state price p^* . The conditions $\mathbb{F}(\bar{z}, \theta; p^*) = 0$ and $\mathbb{H}(\bar{z}, \theta) = 0$ represent match separation and match creation, respectively. The threshold \bar{z}_{\min} solves Equation (19). Finally, c^* , p^* , \bar{z}^* , and θ^* denote the steady-state levels of consumption, price, reservation transportation cost, and product market tightness, respectively.

3.4. Comparative Statics

We use comparative statics to examine how the aggregates in our model respond to unanticipated adverse shocks to aggregate demand, productive capacity, and the supply chain when the economy is at the steady state. These responses provide the identification restrictions for studying the causal effects of supply chain disturbances in the SVAR model in the next section. To address the indeterminacy that arises under a supply chain shock, we assume that the reservation transportation cost is rigid in the short run (fixed by freight contracts), an assumption previewed in Section 2.5 and justified in greater detail later in this section. Appendix E.5 complements the analysis by presenting the full transition dynamics of the model after each shock and showing that they are numerically consistent with the comparative statics discussed below.

An adverse aggregate demand shock arises from either a reduction in the nominal money supply, μ , or a decline in households' preference for consumption, χ . An adverse productive capacity shock corresponds to a negative disturbance to producers' fixed-factor endowment, l . An adverse supply chain shock is modeled as a general increase in transportation costs, represented

by a rise in γ , the scale parameter of the log-normal distribution of transportation costs $G(\cdot)$.

Table 1 summarizes the signs of the responses of macroeconomic aggregates to each shock of interest.²⁴ Figure 6 illustrates the comparative statics of aggregate demand, productive capacity, and supply chain shocks in the downstream retailer–household market (left panels) and the upstream producer–retailer market (right panels).

Table 1: Comparative Statics for Adverse Shocks to Aggregate Demand, Productive Capacity, and the Supply Chain

Adverse Shock To:	Effects On:					
	Consumption (or Output)	Price	Reservation Transportation Cost	Product Market Tightness	Wholesale Price	Spare Capacity
	c	p	\bar{z}	θ	r	$l - c$
Aggregate Demand ($\mu \downarrow$ or $\chi \downarrow$)	–	–	–	–	–	+
Productive Capacity ($l \downarrow$)	–	+	+	+	+	–
<i>Assuming sticky \bar{z} in short-run,</i>						
Supply Chain ($\gamma \uparrow$)	–	+	0	–	±	+

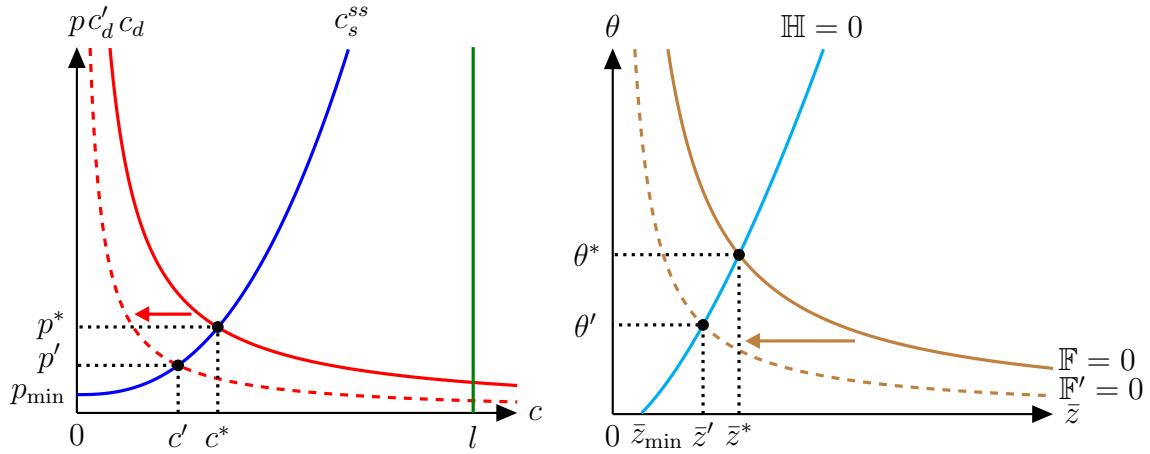
Notes. “0” denotes unchanged, “+” an increase, “–” a decrease, and “±” an undetermined effect. See Appendix E.4 for the derivations of the comparative statics for each shock of interest.

Aggregate demand shock. Panel 6a in Figure 6 illustrates the comparative statics of a decline in aggregate demand, where the demand curve shifts inward from c_d to c'_d as households either hold less money or choose to reduce their consumption of goods.²⁵ The resulting adjustment is a decline in the price required to clear the downstream retailer–household market.

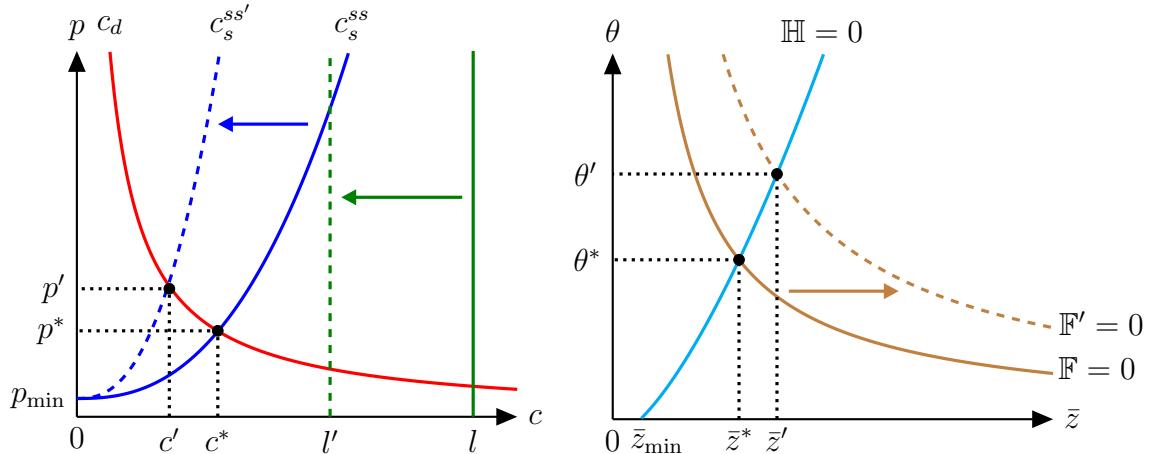
As prices fall and retailers’ profits from household sales decline, retailers visit fewer producers, thereby reducing product market tightness. At the same time, the lower price reduces the reservation transportation cost, since trading partners can only sustain lower costs for trades to remain profitable. The joint decline in price and tightness also lowers the wholesale price: the sale of goods becomes less profitable, and the likelihood of retailers establishing a match with producers increases. Consequently, both reduced product market tightness and lower reservation transportation costs induce producers to sell a smaller share of their productive capacity in the upstream producer–retailer market. Downstream consumption (equivalently, output) falls, and

²⁴Comparable identification restrictions for prices and output can also be derived from a New Keynesian model, where nominal rigidities —rather than search frictions— constrain price adjustment.

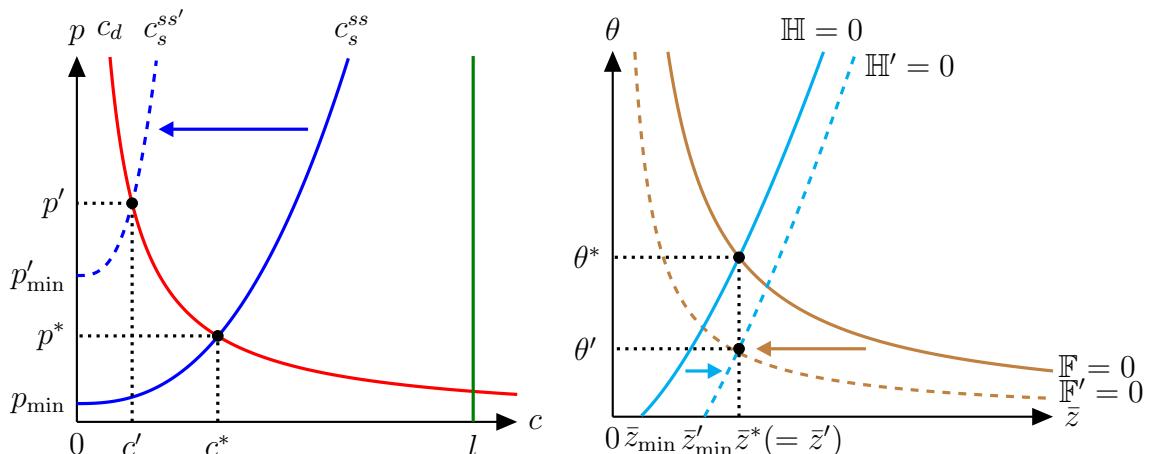
²⁵If a government were present, the same reasoning would apply to fiscal shocks. Hence, our demand shocks capture fiscal and monetary policy shocks, as well as preference shocks.



(a) An Adverse Shock to Aggregate Demand ($\mu \downarrow$ or $\chi \downarrow$)



(b) An Adverse Shock to Productive Capacity ($l \downarrow$)



(c) An Adverse Shock to the Supply Chain ($\gamma \uparrow$)

Figure 6: Graphical Representation of Comparative Statics

spare productive capacity rises.

Productive capacity shock. Panel 6b in Figure 6 shows the effects of a negative supply shock that reduces productive capacity from l to l' . This shock rotates the aggregate supply curve inward while leaving p_{\min} unchanged, since the distribution of transportation costs—and thus the minimum reservation transportation cost \bar{z}_{\min} and the price required for profitable trades—remains unaffected. To clear the retailer–household market, the price must rise.

The higher price attracts more retailers, increasing product market tightness, and simultaneously raises the reservation transportation cost, allowing matches that would otherwise have been forgone to proceed. However, neither the tighter producer–retailer market nor the higher reservation transportation cost is sufficient to offset the direct loss of productive capacity, so aggregate supply falls and consumption declines. At the same time, the joint increase in prices and product market tightness pushes up wholesale prices, while spare capacity contracts.

Supply chain shock. Panel 6c in Figure 6 shows the comparative statics of an increase in the scale parameter of the log-normal distribution of transportation costs, interpreted as an adverse supply chain shock. A higher γ shifts the transportation-cost distribution to the right, increasing the likelihood that producers draw a transportation cost above the reservation threshold. On the one hand, this lowers the expected total surplus from potential producer–retailer matches, dampening retailers’ incentives to search and leading to a less tight market; graphically, the match creation schedule shifts down from $\mathbb{H} = 0$ to $\mathbb{H}' = 0$. On the other hand, the lower expected surplus also discourages producers from engaging in trade, further slackening the producer–retailer market; graphically, the match separation schedule shifts down from $\mathbb{F} = 0$ to $\mathbb{F}' = 0$.

Consequently, the partial equilibrium in the upstream producer–retailer market features lower product market tightness and an *ex ante* ambiguous effect on the reservation transportation cost. Because the response of this reservation cost is undetermined, Equation (17) implies that the resulting impact on aggregate supply—and therefore on the steady-state price in the downstream retailer–household market—remains unclear.

To resolve this analytical indeterminacy, we rely on sectoral evidence of short-run rigidity in containerized shipping, discussed in Section 2.5, and assume that the reservation transportation cost is fixed by freight contracts in the short run. Short-run rigidities in containerized shipping are critical for the on-impact effects of supply chain shocks, since such disturbances operate directly through transportation costs that are typically governed by contractual arrangements, while trade

volumes and profitability do not adjust immediately. By contrast, demand and capacity shocks primarily affect trade volumes or productive potential, which influence firms' willingness to pay for transportation and thereby allow reservation costs to adjust more flexibly.²⁶

With upstream production slack and a sluggish reservation transportation cost, the fraction of goods shipped from producers falls, reducing the supply of goods available in the downstream retailer–household market and pushing up the market-clearing price. Thus, prices rise while consumption falls. Graphically, the aggregate supply curve shifts inward from c_s^{ss} to $c_s^{ss'}$, accompanied by an increase in p_{\min} (the price threshold needed to “turn on” supply). Since the economy's productive capacity is unchanged while the number of successful trades declines, spare capacity rises.²⁷

Finally, the higher downstream retail price feeds back into the upstream production market by increasing the profitability of trades and thereby tightening the market, given that the reservation transportation cost cannot adjust to absorb the additional pressure of the higher price on match separation. To prevent the increase in prices from being strong enough to overturn the immediate slackening in the upstream market caused by the supply chain disturbance (graphically represented by an upward shift of the match separation schedule in Panel 6c), we impose the following bound on the derivative of product market tightness θ with respect to the cost parameter γ :

$$\theta_\gamma \in \left[-\frac{\mathbb{F}_\gamma}{\mathbb{F}_\theta}, 0 \right),$$

where \mathbb{F}_γ and \mathbb{F}_θ denote the partial derivatives of the match separation condition $\mathbb{F}(p, \bar{z}, \theta; \gamma) = 0$ with respect to γ and θ , respectively. Appendix E.4 provides a detailed derivation of this boundary condition. In words, it limits the extent of upstream slack so that the price-feedback effect on match separation does not outweigh the direct effect of the supply chain shock, ensuring the coexistence of a fall in upstream market tightness and a rise in the downstream retail price.

4. The Causal Effects of Supply Chain Disruptions

We now turn to the causal effects of supply chain disruptions using an SVAR model that incorporates our ACR index and imposes restrictions on the responses of macroeconomic aggregates

²⁶Appendix E.5 shows that, even without imposing rigidity, the dynamic responses of all variables —except the reservation transportation cost itself— remain consistent with their analytical counterparts.

²⁷From Equation (9), as prices increase and product market tightness falls in response to a supply chain shock, the effect on the wholesale price is indeterminate.

to three distinct shocks, in line with the theoretical predictions in Table 1.²⁸

4.1. Empirical Framework

Our empirical specification of the SVAR model follows Rubio-Ramírez et al. (2010) and Arias et al. (2018):

$$\mathbf{y}'_t \mathbf{A}_0 = \sum_{l=1}^L \mathbf{y}'_{t-l} \mathbf{A}_l + \boldsymbol{\omega}'_t \mathbf{C} + \boldsymbol{\epsilon}'_t, \quad 1 \leq t \leq T, \quad (21)$$

where \mathbf{y}_t is an $n \times 1$ vector of endogenous variables, $\boldsymbol{\omega}_t = [1, t]'$ is a 2×1 vector containing a constant and a linear trend, and $\boldsymbol{\epsilon}_t$ is an $n \times 1$ vector of structural shocks. The matrices \mathbf{A}_l are $n \times n$ structural parameter matrices for $0 \leq l \leq L$, with \mathbf{A}_0 assumed invertible, and \mathbf{C} is a $2 \times n$ parameter matrix. The lag length is denoted by L , and T is the sample size. Conditional on past information and the initial conditions $\mathbf{y}_0, \dots, \mathbf{y}_{1-L}$, the vector $\boldsymbol{\epsilon}_t$ is Gaussian with mean zero and variance-covariance matrix \mathbf{I}_n , the $n \times n$ identity matrix.

Equation (21) can be written more compactly as:

$$\mathbf{y}'_t \mathbf{A}_0 = \mathbf{x}'_t \mathbf{A}_+ + \boldsymbol{\epsilon}'_t, \quad 1 \leq t \leq T, \quad (22)$$

where $\mathbf{A}'_+ = [\mathbf{A}'_1 \dots \mathbf{A}'_L \mathbf{C}']$ and $\mathbf{x}'_t = [\mathbf{y}'_{t-1} \dots \mathbf{y}'_{t-L} \boldsymbol{\omega}'_t]$. The dimension of \mathbf{A}_+ is $m \times n$, where $m = nL + 2$. The reduced-form representation implied by Equation (22) is:

$$\mathbf{y}'_t = \mathbf{x}'_t \mathbf{B} + \mathbf{u}'_t, \quad 1 \leq t \leq T,$$

where $\mathbf{B} = \mathbf{A}_+ \mathbf{A}_0^{-1}$, $\mathbf{u}'_t = \boldsymbol{\epsilon}'_t \mathbf{A}_0^{-1}$, and $\mathbb{E}(\mathbf{u}_t \mathbf{u}'_t) = \boldsymbol{\Sigma} = (\mathbf{A}_0 \mathbf{A}'_0)^{-1}$.

Motivated by the variables in our theoretical model, we estimate the baseline SVAR using monthly U.S. data on real personal consumption expenditures (PCE), the PCE chain-type price index, and the import price index for all commodities. We also include two empirical measures constructed from Equations (23) and (24) —spare productive capacity and product market tightness— together with our ACR index. The sample spans January 2016 to March 2025, and all series are seasonally adjusted. Appendix F provides a detailed overview of the external data and their sources.

²⁸We could also pursue a full structural estimation of our model. However, this would require numerous ancillary assumptions (e.g., parametric forms, shock persistence) that may be unreliable given current knowledge of global supply chain models. Although we rely on some of these assumptions for our identification restrictions, we are cautiously optimistic that these restrictions hold for more general specifications, even if demonstrable only numerically. Hence, the added flexibility of SVARs appears most suitable here.

As mentioned above, we first construct an empirical measure of spare productive capacity ($\text{SpareCapacityRate}_t$) by computing the import-weighted average spare capacity rate of the top five U.S. trading partners (Mexico, Canada, China, Germany, and Japan), which together account for more than half of U.S. goods imports.²⁹ The weights are based on U.S. goods imports from each country in 2016. Formally,

$$\text{SpareCapacityRate}_t = \sum_{i \in \mathcal{C}} \left[\frac{\text{Import}_{i,2016}}{\sum_{i \in \mathcal{C}} \text{Import}_{i,2016}} \cdot (1 - \text{CapacityUtilization}_{i,t}) \right], \quad (23)$$

where $\mathcal{C} \equiv \{\text{Mexico, Canada, China, Germany, Japan}\}$, $\text{Import}_{i,2016}$ denotes U.S. goods imports (customs basis) from country i in 2016, and $\text{CapacityUtilization}_{i,t}$ is the capacity utilization rate for country i in month t .

Next, we construct an empirical measure of product market tightness (Tightness_t). Following the definition $\theta \equiv i_U/x_U$, we proxy the number of unmatched retailers (i_U) with total U.S. manufacturers' new orders and the number of unmatched producers (x_U) with the import-weighted average spare capacity of the specified U.S. trading partners. Specifically,

$$\text{Tightness}_t = \frac{\text{ManufactureNewOrder}_t}{\text{SpareCapacityDollar}_t}. \quad (24)$$

The time series for U.S. manufacturers' new orders ($\text{ManufactureNewOrder}_t$) is from the Federal Reserve Economic Data (FRED). We construct the U.S. dollar value of the import-weighted average spare capacity ($\text{SpareCapacityDollar}_t$) as follows:

$$\text{SpareCapacityDollar}_t = \sum_{i \in \mathcal{C}} \left[\frac{\text{Import}_{i,2016}}{\sum_{i \in \mathcal{C}} \text{Import}_{i,2016}} \cdot \left(\frac{\text{IP}_{i,t}}{\text{CapacityUtilization}_{i,t}} - \text{IP}_{i,t} \right) \right], \quad (25)$$

where $\text{IP}_{i,t}$ denotes total industrial production of country i in month t , measured in millions of constant 2005 U.S. dollars. Finally, we use the import price as a proxy for the wholesale price to capture U.S. manufacturers' and retailers' international sourcing strategies, particularly during the COVID-19 pandemic.

Real PCE, the PCE price index, product market tightness, and the import price index enter the SVAR in log points, while spare capacity and the ACR index enter in percentages. In the baseline specification, we set the lag length to two, although we check that the results are robust to longer lags.

²⁹See <https://wits.worldbank.org/CountrySnapshot/en/USA> (accessed October 15, 2024).

We estimate the SVAR model using the Bayesian approach of [Arias et al. \(2018, 2019, 2023\)](#), imposing restrictions only on the first period of impulse responses (horizon $k = 1$), thereby following the minimal-structure strategy of [Mumtaz and Zanetti \(2012, 2015\)](#).³⁰ As shown in Appendix [E.5](#), the dynamic version of our theoretical model converges quickly —within one month— and monotonically from one steady state to another following each shock of interest. This provides additional support for imposing identification restrictions only at horizon $k = 1$.

Our identification scheme applies the sign restrictions derived from our theoretical model, summarized in Table [1](#), together with zero restrictions on the response of the ACR index to adverse shocks to aggregate demand and productive capacity at $k = 1$:

Restriction 1. *An adverse shock to aggregate demand leads to a negative response of real PCE, the PCE price index, product market tightness, and the import price index, as well as a positive response of spare capacity at $k = 1$. The ACR index does not respond at $k = 1$.*

Restriction 2. *An adverse shock to productive capacity leads to a negative response of real PCE and spare capacity, and a positive response of the PCE price index, product market tightness, and the import price index at $k = 1$. The ACR index does not respond at $k = 1$.*

Restriction 3. *An adverse shock to the supply chain leads to a negative response of real PCE and product market tightness, and a positive response of the PCE price index, spare capacity, and the ACR index at $k = 1$.*

The zero restrictions on the ACR index in Restrictions [1](#) and [2](#) are motivated by the short-run operational rigidity of the containerized shipping industry, as discussed in Section [2.5](#). Container ships typically require several weeks' notice to alter schedules and arrivals in response to demand or capacity shocks. As a result, our ACR index should not react within the first month. This is a conservative assumption, but, as shown later in this section, it enables sharper identification (i.e., narrower posterior probability bands) of the inflationary nature of supply chain shocks compared with cases where such zero restrictions cannot be imposed.

³⁰We adopt a Normal-Generalized-Normal (NGN) prior distribution over \mathbf{A}_0 and \mathbf{A}_+ . The NGN prior is a conjugate prior characterized by four parameters $(\nu, \Phi, \Psi, \Omega)$. Parameters ν and Φ govern the marginal prior distribution of $\text{vec}(\mathbf{A}_0)$, while Ψ and Ω govern the prior distribution of $\text{vec}(\mathbf{A}_+)$ conditional on \mathbf{A}_0 . We set $\nu = 0$, $\Phi = \mathbf{0}_{n \times n}$, $\Psi = \mathbf{0}_{m \times n}$, and $\Omega^{-1} = \mathbf{0}_{m \times m}$. This parameterization yields prior densities equivalent to those in [Uhlig \(2005\)](#). Appendix [G.1](#) further shows that our results are robust to the prior-robust approach of [Giacomini and Kitagawa \(2021\)](#).

By contrast, Restriction 3 allows the ACR index to respond to a supply chain disturbance at $k = 1$, though it does not require such a response —we assess whether the corresponding posterior probability bands exclude zero.

4.2. Baseline Results

Figures 7, 8, and 9 present the responses of the endogenous variables to adverse shocks to aggregate demand, productive capacity, and the supply chain, respectively. The solid lines depict the point-wise posterior median impulse response functions (IRFs), and the shaded areas show the 68% and 90% posterior probability bands. Unless otherwise noted, all results are based on ten million orthogonal reduced-form draws $(\mathbf{B}, \Sigma, \mathbf{Q})$, where \mathbf{Q} is an $n \times n$ orthogonal matrix.³¹

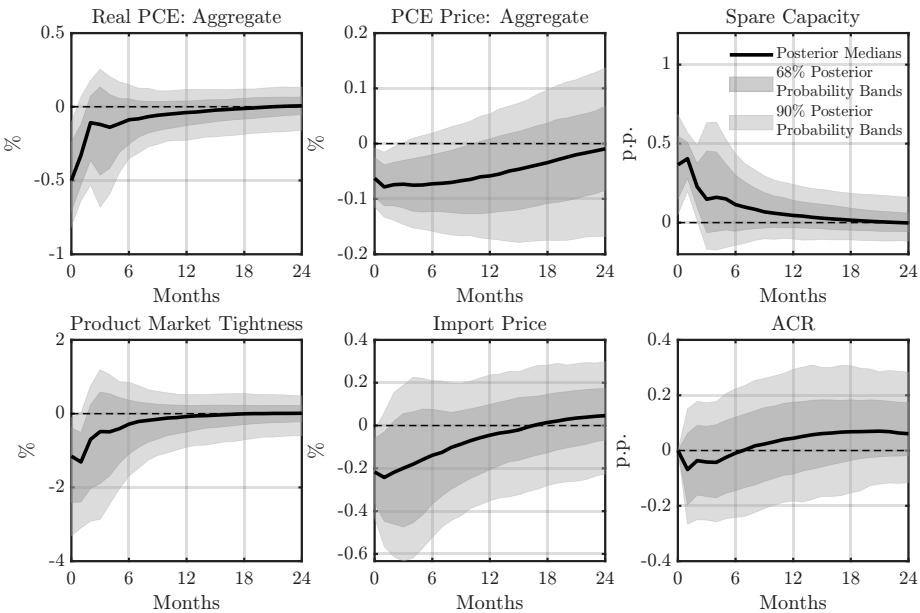


Figure 7: IRFs to an Adverse Shock to Aggregate Demand

Notes. The IRFs to a one-standard-deviation adverse aggregate demand shock are identified using the Bayesian SVAR in Equation (21), with the ACR index included and the sign/zero restrictions at horizon $k = 1$ imposed as specified in Restrictions 1-3.

We begin with the IRFs to an adverse aggregate demand shock, as shown in Figure 7. On impact, real PCE declines by about 0.5%, while spare capacity increases by slightly less than 0.5 percentage point (p.p.). Both responses are precisely estimated, with posterior probability bands excluding zero for roughly three months. Product market tightness exhibits a sharp drop of more than 1%, and the lower bound of the 90% probability band suggests the effect could be

³¹Two representations $(\mathbf{A}_0, \mathbf{A}_+)$ and $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)$ are observationally equivalent if and only if $\mathbf{A}_0 = \tilde{\mathbf{A}}_0 \mathbf{Q}$ and $\mathbf{A}_+ = \tilde{\mathbf{A}}_+ \mathbf{Q}$ for some orthogonal \mathbf{Q} (Arias et al., 2018).

as large as 3–4%; tightness then gradually converges back toward zero within two years. The PCE price index decreases modestly by about 0.1% on impact, remaining negative with high posterior probability for more than a year before fading. The import price index displays a similar deflationary pattern, initially falling by 0.2% and remaining significantly below baseline through the first year. Lastly, the ACR index shows a short-lived negative response, but the wide probability bands indicate high uncertainty, with much of the posterior mass centered near zero.

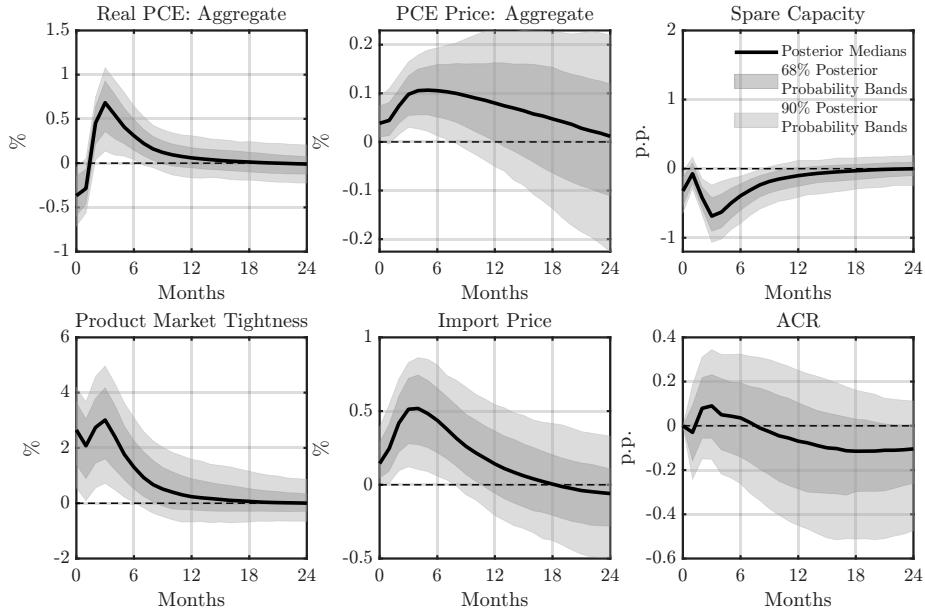


Figure 8: IRFs to an Adverse Shock to Productive Capacity

Notes. The IRFs to a one-standard-deviation adverse productive capacity shock are identified using the Bayesian SVAR in Equation (21), with the ACR index included and the sign/zero restrictions at horizon $k = 1$ imposed as specified in Restrictions 1–3.

Figure 8 reports the IRFs to an adverse productive capacity shock. On impact, real PCE dips slightly and spare capacity falls, while the PCE price, the import price, and product market tightness rise, in line with Restriction 2. Tightness peaks on impact (around 2–3%), dips briefly, and then rises again at roughly the three-month mark before gradually trending back toward zero over the following year (an impulse pattern that, when interpreted through our theoretical model, reflects a sudden contraction in productive capacity that raises tightness immediately, recall Panel 6b in Figure 6, with slower price pass-through subsequently drawing additional unmatched retailers and producing a secondary hump).

Import prices increase by about 0.5% within one quarter and remain above baseline for six quarters, whereas PCE prices rise more modestly (peaking near 0.1%) and fade more slowly. Real

PCE reverses quickly from its initial dip, rising by about 0.5-0.6% within the first quarter before mean-reverting; by contrast, spare capacity troughs just above -0.5 p.p. around the one-quarter mark and then recovers gradually toward zero. The ACR index is muted on impact, climbs to roughly 0.1 p.p. at about three months, and then drifts below zero around the three-quarter horizon, stabilizing near -0.1 p.p., suggesting complementarities between production constraints and transportation frictions in the transmission of capacity shocks.

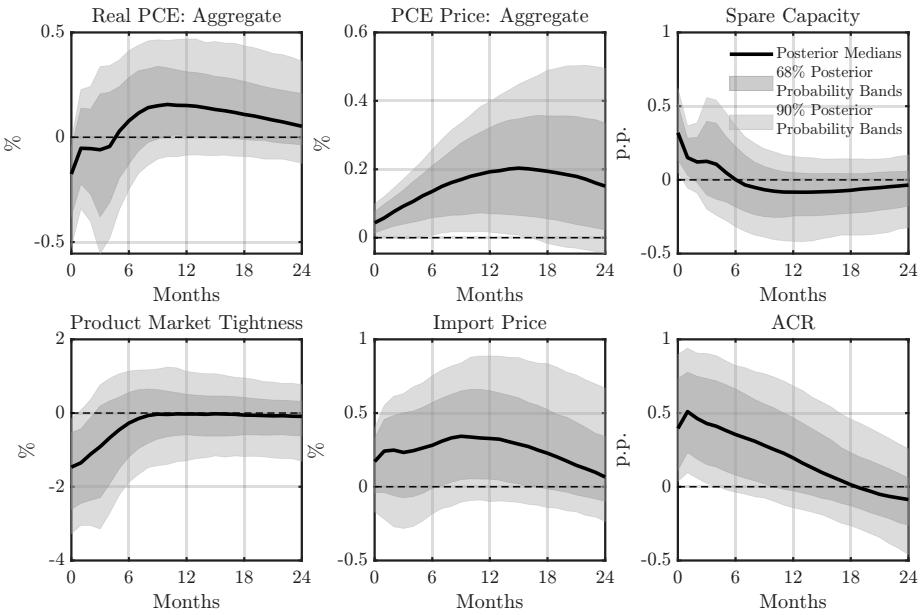


Figure 9: IRFs to an Adverse Shock to the Supply Chain

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are identified using the Bayesian SVAR in Equation (21), with the ACR index included and the sign/zero restrictions at horizon $k = 1$ imposed as specified in Restrictions 1-3.

Figure 9 shows the IRFs to an adverse supply chain shock. On impact, real PCE and product market tightness fall, while spare capacity and the ACR index rise, in line with our identification scheme. The PCE price response is sizable: the median increases steadily, peaking near 0.2% around one year, and the lower bound of the 90% probability band stays above zero throughout, indicating a robust inflationary effect.

Import prices, which are unrestricted in our specification, also rise persistently. The 68% probability band remains above zero across most horizons, and the median peaks above 0.3% at roughly the three-quarter mark. In the decomposition implied by Equation (9), this persistence reflects that the direct price component outweighs the slackening (i.e., lower tightness) component, yielding a net positive effect on import prices.

The real-side variables display a distinct dynamic pattern. Real PCE dips modestly on impact (about -0.2%), turns positive within 4-6 months (peaking near 0.2%), and then gradually mean-reverts. Spare capacity rises initially (0.3 p.p.) but only briefly: it crosses below zero after a few months and remains slightly negative for about a year before returning toward baseline. The ACR index jumps on impact (0.5 p.p.) and then decays gradually, crossing zero around the 18-month horizon, consistent with congestion easing as logistics adjust.

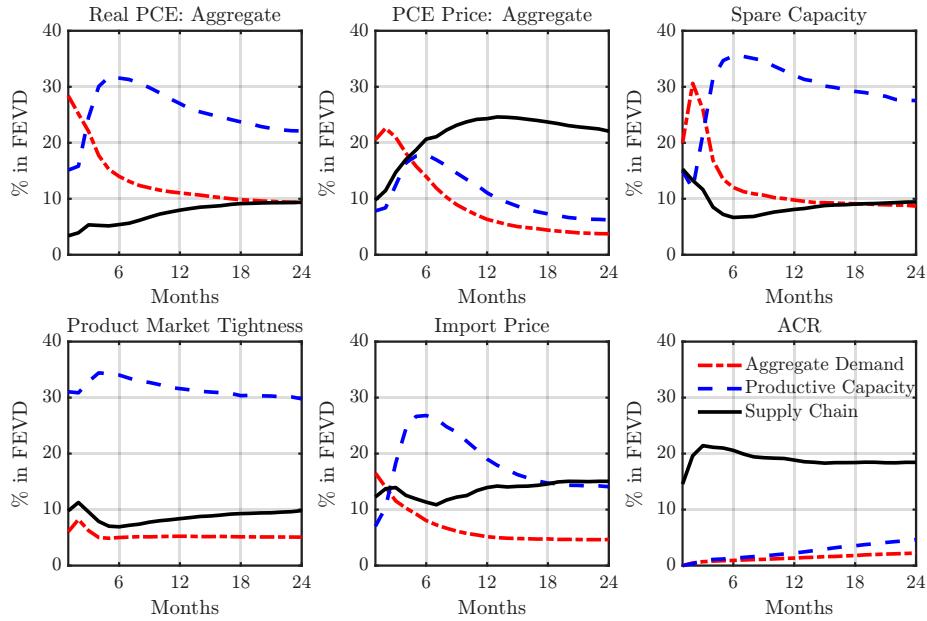


Figure 10: Forecast Error Variance Decompositions (FEVDs) from the SVAR

Notes. Each line shows the posterior median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the identified Bayesian SVAR in Equation (21), with the ACR index included and the sign/zero restrictions imposed at horizon $k = 1$ (Restrictions 1-3).

Figure 10 shows the share of forecast error variance explained by each of the three structural shocks. Two broad patterns emerge. First, price variables are mainly driven by supply chain disturbances at medium horizons. For the PCE price index, the supply chain share rises from roughly 10% at short horizons to about 20-25% by 6-12 months and remains the largest contributor thereafter, while the contributions of demand and capacity shocks decline toward 5-10%. For import prices, productive capacity shocks lead early (peaking near 25% around six months), but the supply chain share increases steadily and eventually overtakes the capacity share, reaching roughly 15% by two years; demand remains comparatively small throughout.

Second, real activity and product market tightness are driven primarily by productive capacity

shocks. For real PCE, capacity shocks account for roughly 30% at 3-6 months and remain the largest contributor even at two years, while the demand share is front-loaded —high on impact but falling toward 10%— and the supply chain share rises gradually into the high single digits. For spare capacity and product market tightness, capacity shocks again dominate (around 30-35% at their peaks), with demand contributing mainly near impact and the supply chain share remaining modest, though drifting upward over time.

As a validation check, the ACR is explained primarily by supply chain shocks —about 15-21% across horizons— while demand and capacity shocks contribute little, especially within the first six months. This aligns with the view that port congestion is largely insulated from short-run fluctuations in demand and productive capacity, as discussed in Section 2.5.³²

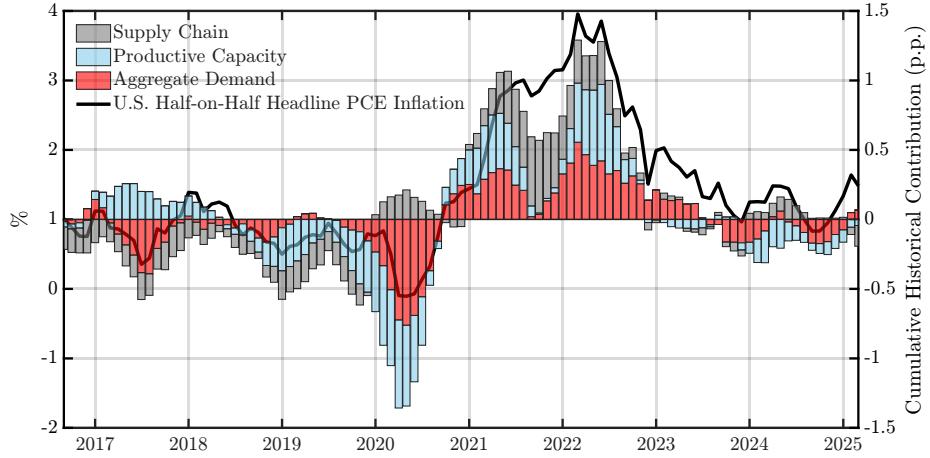


Figure 11: Historical Decomposition (HD) of U.S. HoH Headline PCE Inflation

Notes. The solid line plots U.S. headline PCE inflation, measured as HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars plot the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The historical decomposition is computed from the identified Bayesian SVAR in Equation (21), using first-differenced endogenous variables, and we report posterior medians. Sign/zero restrictions are imposed at horizon $k = 1$ as specified in Restrictions 1-3.

Figure 11 presents the key empirical result of our analysis: the cumulative historical contributions of the three identified shocks to U.S. half-on-half (HoH) headline PCE inflation over the period September 2016 to March 2025.³³ The main findings from the decomposition can be

³²Shares sum to less than 100% because only the three identified shocks are plotted, with the residual variance attributed to unidentified shocks.

³³Historical contributions are accumulated beginning in September 2016, reflecting two lags in the VAR and the use of first-differenced variables to compute HoH growth rates. For clarity, the figure excludes contributions from lag terms, exogenous variables (the constant and the linear trend), and unidentified shocks; consequently, the stacked bars for the identified shocks do not exactly sum to the black line (U.S. HoH headline PCE inflation).

summarized in five points.

1. Before the pandemic (2016-2019), contributions from all three shocks were modest and often offsetting, with the supply chain component generally slightly negative.³⁴
2. The sharp disinflation at the onset of COVID-19 (early 2020) was driven primarily by a large negative aggregate demand contribution, plausibly reflecting mobility restrictions that suppressed consumption and heightened uncertainty, and reinforced by a concurrent negative contribution from productive capacity, likely dominated by the sharp plunge in global oil prices.
3. The inflation surge from 2021 through mid-2022 was dominated by supply chain shocks, whose cumulative contribution rose rapidly and surpassed that of the other shocks. Aggregate demand also turned positive, and productive capacity shocks became increasingly inflationary due to the gradual increase in global oil prices, China's Omicron-related mobility restrictions, and labor-market frictions in major U.S. trading partners that constrained effective capacity.
4. From the second half of 2022 through 2023, disinflation reflected the unwinding of supply chain pressures and improving capacity. Demand faded as aggressive interest rate hikes and quantitative tightening by major central banks worked to counter the high inflation inherited from 2022.
5. From early 2024 to the end of the sample, inflation reaccelerated mildly, largely because of renewed supply chain pressures (e.g., the escalation of the Red Sea crisis; see [Bai et al., 2025a](#) for details), while demand and capacity continued to exert a modest drag.³⁵

4.3. Robustness Checks and Extensions

We conduct a range of robustness checks on our baseline results beyond those previously discussed. First, [Appendix G.2](#) re-estimates the SVAR using monthly inflation instead of price levels for PCE and import prices. The conclusions are unchanged: cumulated inflation IRFs closely mirror the price-level IRFs, and FEVD rankings remain stable. In addition, the shape

³⁴This finding highlights the importance of strategic supply chain improvements to mitigate inflationary pressures. For instance, U.S. ports, such as the Port of Los Angeles, undertook infrastructure upgrades between 2017 and 2019 to enhance capacity, efficiency, and resilience against systemic disruptions.

³⁵Two additional observations are worth noting: (i) the peak inflation period aligns with the largest positive supply chain contribution, consistent with the strong price effects in [Figure 9](#); and (ii) during the subsequent disinflation, capacity plays the leading role on the downside, in line with FEVD patterns that assign more real-side variation to capacity shocks.

and magnitude of the inflation IRFs to demand and capacity shocks align with classic evidence on inflation responses to demand- and supply-side disturbances (e.g., [Smets and Wouters, 2007](#)).

Second, Appendix [G.3](#) examines invertibility by augmenting the information set in the SVAR. Adding services activity and prices, oil prices, the federal funds rate, and wages leaves demand- and capacity-driven IRFs essentially unchanged, but attenuates the supply chain price effects, as one would expect from having more endogenous variables.

Third, Appendix [G.4](#) tests for nonlinear transmission by shock size. Using LPs that interact the identified shock series with indicators for small versus large realizations, and integrating over the full posterior distribution of shocks from the Bayesian SVAR, we find divergent posterior median IRFs but overlapping credible bands. This pattern suggests possible nonlinearities — particularly for large shocks (e.g., [Baquee and Farhi, 2019, 2020](#)), yet our main conclusions are robust: the transmission mechanism is qualitatively similar across shock sizes, and the macroeconomic effects of the three identified shocks shown in Figures [7-9](#) are not sensitive to moderate nonlinear amplification.

Fourth, Appendix [G.5](#) shows that re-estimating the monthly SVAR with four and six lags — holding priors, identification, and sample fixed — leaves the median IRFs and qualitative dynamics essentially unchanged across all three shocks, with only modest widening of posterior probability bands at short horizons. Finally, Appendices [G.6](#) and [G.7](#) provide further checks by dropping the linear trend and by using alternative proxies for activity and prices.

We also explore two extensions. In a companion paper ([Bai et al., 2025b](#)), we analyze the effectiveness of monetary policy during global supply chain disruptions. Guided by the theoretical predictions of our model in Section [3](#) and supported by evidence from a threshold vector autoregression (TVAR) and LPs, we find that supply chain disruptions increase the sensitivity of prices while reducing the sensitivity of output to contractionary monetary policy shocks, thereby introducing state dependence into the stabilization trade-off for monetary policy.

Furthermore, we re-estimate the U.S. SVAR with goods and services inflation and three shocks (demand reallocation, oil, and supply chain). The pandemic-era surge in goods inflation is driven mainly by demand reallocation (e.g., [Guerrieri et al., 2022](#); [Ferrante et al., 2023](#)) and supply chain shocks, with oil playing a visible but secondary role. By contrast, services inflation moves little, shows limited pass-through from goods, and reflects only modest contributions from the identified shocks, consistent with stickier prices and sector-specific dynamics. For the Euro Area,

an analogous SVAR points to a dominant and persistent role for oil price shocks in 2021-22, reflecting Europe's heightened energy dependence. Full results are available upon request.

4.4. Identification Gains From Congestion Indices

Finally, we directly compare the estimated price responses across alternative proxies for global supply chain conditions, illustrating the identification gains from using our port congestion indices to isolate supply chain disturbances.

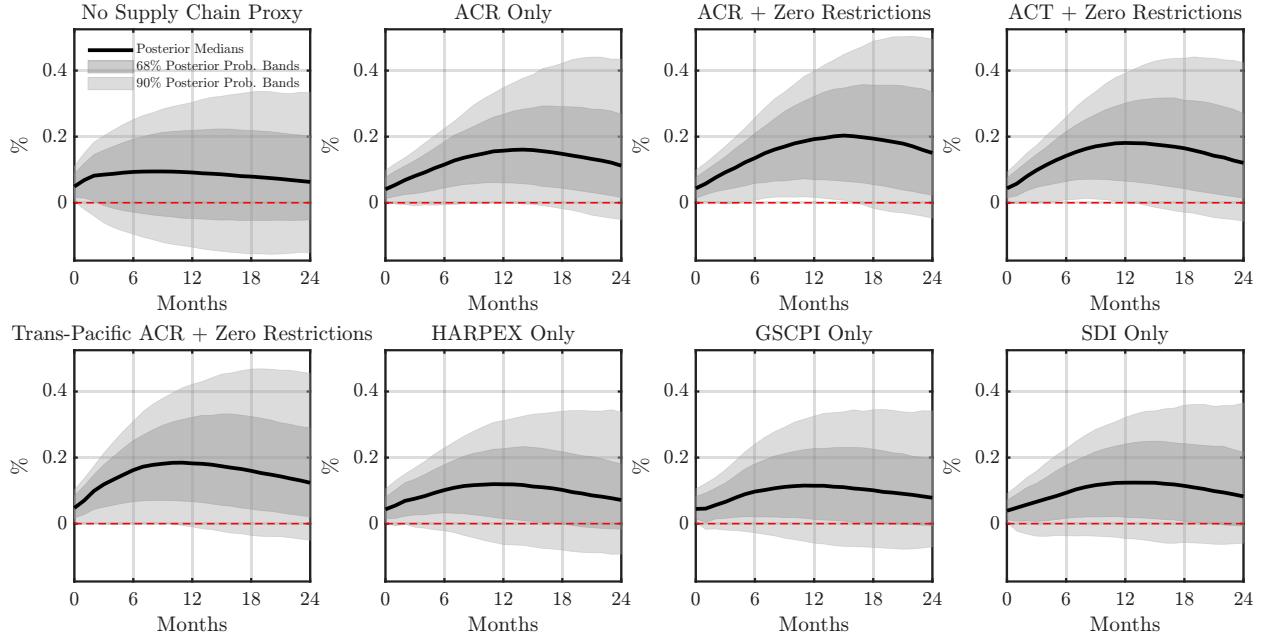


Figure 12: Price Responses to an Adverse Supply Chain Shock Across Alternative Proxies

Notes. IRFs of the U.S. PCE chain-type price index to a one-standard-deviation adverse supply chain shock, identified using the Bayesian SVAR in Equation (21) across eight specifications: (i) no supply chain proxy; (ii) ACR; (iii) ACR with zero restrictions at horizon $k = 1$ on the responses of ACR to aggregate demand and productive capacity shocks; (iv) ACT with the same zero restrictions; (v) Trans-Pacific ACR with the same zero restrictions; (vi) HARPEX; (vii) the New York Fed's GSCPI (Benigno et al., 2022); and (viii) the SDI from Smirnyagin and Tsyvinski (2022) and Liu et al. (2024). The sign restrictions in Restrictions 1-3 are imposed to identify the adverse supply chain shock in all specifications except (i), where the positive restriction on the ACR response to such a shock is omitted.

Figure 12 shows that, when the SVAR is augmented with the ACR index, the IRF of the PCE chain-type price index to an adverse supply chain shock is larger in magnitude and more precisely estimated than when no proxy is included. Imposing zero restrictions on the responses of the ACR (ACT, or the Trans-Pacific ACR, i.e., a targeted ACR index for major ports along the Trans-Pacific route) to demand and capacity shocks at horizon $k = 1$ —motivated by our domain knowledge of short-run operational rigidity in the containerized shipping industry, as

discussed in Section 2.5—further sharpens identification. The resulting price responses exhibit higher peaks, and the lower bound of the 90% probability band diverges from the zero-response line as the median response approaches its peak, thereby underscoring the inflationary effects of supply chain disturbances in line with our theoretical predictions. By contrast, when using HARPEX, GSCPI, or SDI, the median responses are smaller and the credible bands straddle zero at all horizons, indicating weaker and less precisely estimated price effects.

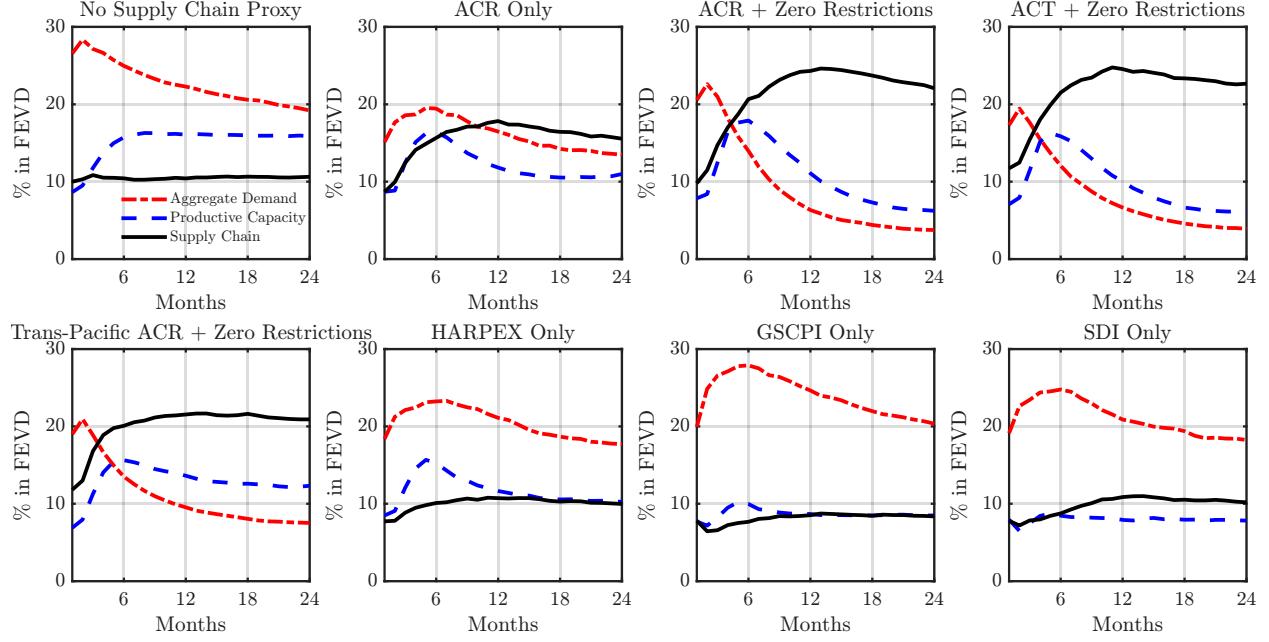


Figure 13: FEVD of the U.S. PCE Price Index Across Alternative Proxies

Notes. The figure plots the posterior-median shares of the forecast error variance of the U.S. PCE chain-type price index attributable to aggregate demand, productive capacity, and supply chain shocks across horizons, for each of the eight specifications discussed above.

Consistent with these IRFs, the FEVD results in Figure 13 show that using our congestion indices—especially under the above zero restrictions—attributes a larger share of the forecast error variance of PCE prices to supply chain shocks at medium horizons. By contrast, when alternative proxies are employed, the variance attribution shifts toward aggregate demand, with supply chain shocks playing only a minimal role throughout. We report the full IRF and FEVD comparisons, along with a robustness check that replaces the PCE chain-type price index with the PCE goods price index (which yields analogous conclusions), in Appendix H.

5. Conclusion

Our findings suggest several promising avenues for future research. First, the ACR index reveals substantial heterogeneity in global port congestion, raising questions about whether spillovers across ports stem from geographical proximity or production synergies (Fernández-Villaverde et al., 2021, 2024a,b). Second, our results highlight the importance of spare capacity in assessing the severity of supply chain disruptions, suggesting the need to endogenize its adjustment and examine its persistence and interaction with inventories. Third, incorporating input–output networks to analyze how spare capacity transmits supply chain shocks across firms could yield valuable insights, as production structures may amplify or dampen disturbances and trigger endogenous changes (Ghassibe, 2024; Xu et al., 2025). We plan to explore some of these extensions in future work.

References

Acharya, V. V., Crosignani, M., Eisert, T., and Eufinger, C. (2023). How Do Supply Shocks to Inflation Generalize? Evidence From the Pandemic Era in Europe. Working Paper 31790, National Bureau of Economic Research.

Alessandria, G., Khan, S. Y., Khederlarian, A., Mix, C., and Ruhl, K. J. (2023). The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020–2022. *Journal of International Economics*, 146:103788.

Allen, T. and Arkolakis, C. (2014). Trade and the Topography of the Spatial Economy. *Quarterly Journal of Economics*, 129(3):1085–1140.

Antràs, P. (2023). An ‘Austrian’ Model of Global Value Chains. Working Paper 30901, National Bureau of Economic Research.

Arias, J. E., Caldara, D., and Rubio-Ramírez, J. F. (2019). The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure. *Journal of Monetary Economics*, 101:1–13.

Arias, J. E., Fernández-Villaverde, J., Rubio-Ramírez, J. F., and Shin, M. (2023). The Causal Effects of Lockdown Policies on Health and Macroeconomic Outcomes. *American Economic Journal: Macroeconomics*, 15(3):287–319.

Arias, J. E., Rubio-Ramírez, J. F., and Waggoner, D. F. (2018). Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications. *Econometrica*, 86:685–720.

Ascari, G., Bonam, D., and Smadu, A. (2024). Global Supply Chain Pressures, Inflation, and Implications for Monetary Policy. *Journal of International Money and Finance*, 142:103029.

Bai, X., Fernández-Villaverde, J., Li, Y., Marto, R., and Zanetti, F. (2025a). Shipping to America. Working paper, University of Oxford.

Bai, X., Fernández-Villaverde, J., Li, Y., and Zanetti, F. (2025b). State Dependence of Monetary Policy During Global Supply Chain Disruptions. Working paper, University of Pennsylvania.

Bai, X. and Li, Y. (2022). The Congestion Effect of Oil Transportation and Its Trade Implications. Working paper, SSRN.

Bai, X., Ma, Z., Hou, Y., Li, Y., and Yang, D. (2023). A Data-Driven Iterative Multi-Attribute Clustering Algorithm and Its Application in Port Congestion Estimation. *IEEE Transactions on Intelligent Transportation Systems*, 24:12026–12037.

Baqae, D. and Farhi, E. (2019). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem. *Econometrica*, 87(4):1155–1203.

Baqae, D. and Farhi, E. (2020). Nonlinear Production Networks With an Application to the COVID-19 Crisis. Working Paper 27281, National Bureau of Economic Research.

Barro, R. J. and Grossman, H. I. (1971). A General Disequilibrium Model of Income and Employment. *American Economic Review*, 61(1):82–93.

Bauer, M. D. and Swanson, E. T. (2023). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. *NBER Macroeconomics Annual*, 37:87–155.

Benigno, G., di Giovanni, J., Groen, J. J., and Noble, A. I. (2022). The GSCPI: A New Barometer of Global Supply Chain Pressures. Staff Report 1017, Federal Reserve Bank of New York.

Benigno, P. and Eggertsson, G. B. (2023). It’s Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. Working Paper 31197, National Bureau of Economic Research.

Bernanke, B. and Blanchard, O. (2025). What Caused the US Pandemic-Era Inflation? *American Economic Journal: Macroeconomics*, 17(3):1–35.

Binetti, A., Nuzzi, F., and Stantcheva, S. (2024). People’s Understanding of Inflation. *Journal of Monetary Economics*, page 103652.

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2020). Geography, Transportation, and Endogenous Trade Costs. *Econometrica*, 88(2):657–691.

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2024). Investment in Infrastructure and Trade: The Case of Ports. Working Paper 32503, National Bureau of Economic Research.

Brancaccio, G., Kalouptsidi, M., Papageorgiou, T., and Rosaia, N. (2023). Search Frictions and Efficiency in Decentralized Transport Markets. *Quarterly Journal of Economics*, pages 2451–2503.

Bureau of Transportation Statistics (2021). On National Maritime Day and Every Day, U.S. Economy Relies on Waterborne Shipping. Data Spotlights.

Cerdeiro, D. A. and Komaromi, A. (2020). Supply Spillovers During the Pandemic: Evidence From High-Frequency Shipping Data. Working Paper 2020/284, International Monetary Fund.

Chen, L., Zhang, D., Ma, X., Wang, L., Li, S., Wu, Z., and Pan, G. (2016). Container Port Performance Measurement and Comparison Leveraging Ship GPS Traces and Maritime Open Data. *IEEE Transactions on Intelligent Transportation Systems*, 17(5):1227–1242.

Comín, D. A., Johnson, R. C., and Jones, C. J. (2023). Supply Chain Constraints and Inflation. Finance and Economics Discussion Series 2023-075, Federal Reserve Board.

di Giovanni, J., Şebnem Kalemli-Özcan, Silva, A., and Yildirim, M. A. (2022). Global Supply Chain Pressures, International Trade, and Inflation. Working Paper 30240, National Bureau of Economic Research.

di Giovanni, J., Şebnem Kalemli-Özcan, Silva, A., and Yildirim, M. A. (2023). Pandemic-Era Inflation Drivers and Global Spillovers. Working Paper 31887, National Bureau of Economic Research.

Dunn, J. and Leibovici, F. (2023). Navigating the Waves of Global Shipping: Drivers and Aggregate Implications. Working Paper 2023-002, Federal Reserve Bank of St. Louis.

Fan, L., Wilson, W. W., and Dahl, B. (2012). Congestion, Port Expansion and Spatial Competition for US Container Imports. *Transportation Research Part E: Logistics and Transportation Review*, 48(6):1121–1136.

Feng, M., Shaw, S.-L., Peng, G., and Fang, Z. (2020). Time Efficiency Assessment of Ship Movements in Maritime Ports: A Case Study of Two Ports Based on AIS Data. *Journal of Transport Geography*, 86:102741.

Fernández-Villaverde, J., Mandelman, F., Yu, Y., and Zanetti, F. (2021). The “Matthew Effect” and Market Concentration: Search Complementarities and Monopsony Power. *Journal of Monetary Economics*, 121:62–90.

Fernández-Villaverde, J., Mandelman, F., Yu, Y., and Zanetti, F. (2024a). Search Complementarities, Aggregate Fluctuations, and Fiscal Policy. *The Review of Economic Studies*, page rdae053.

Fernández-Villaverde, J., Yu, Y., and Zanetti, F. (2024b). Technological Synergies, Heterogeneous Firms, and Idiosyncratic Volatility. Working Paper 32247, National Bureau of Economic Research.

Ferrante, F., Graves, S., and Iacoviello, M. (2023). The Inflationary Effects of Sectoral Reallocation. *Journal of Monetary Economics*, 140:S64–S81.

Finck, D., Klein, M., and Tillmann, P. (2024). The Inflationary Effects of Global Supply Chain Shocks: Evidence From Swedish Microdata. Working paper, SSRN.

Finck, D. and Tillmann, P. (2022). The Macroeconomic Effects of Global Supply Chain Disruptions. Discussion Paper 14/2022, Bank of Finland Institute for Emerging Economies.

Freightos (2024). Shipping From Shanghai, Shanghai to Los Angeles, CA: Air, Sea & Container Freight. Freightos.

Fuchs, S. and Wong, W. F. (2022). Multimodal Transport Networks. Working Paper 2022-13, Federal Reserve Bank of Atlanta.

Ganapati, S., Wong, W. F., and Ziv, O. (2024). Entrepôt: Hubs, Scale, and Trade Costs. *American Economic Journal: Macroeconomics*, 16(4):239–78.

Ghassibe, M. (2024). Endogenous Production Networks and Non-Linear Monetary Transmission. Working Paper 1449, Barcelona School of Economics.

Ghassibe, M. and Zanetti, F. (2022). State Dependence of Fiscal Multipliers: The Source of Fluctuations Matters. *Journal of Monetary Economics*, 132:1–23.

Giacomini, R. and Kitagawa, T. (2021). Robust Bayesian Inference for Set-Identified Models. *Econometrica*, 89:1519–1556.

Guerrieri, V., Lorenzoni, G., Straub, L., and Werning, I. (2022). Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages? *American Economic Review*, 112(5):1437–74.

Harding, M., Lindé, J., and Trabandt, M. (2023). Understanding Post-COVID Inflation Dynamics. *Journal of Monetary Economics*, 140:S101–S118.

Heiland, I., Moxnes, A., Ulltveit-Moe, K. H., and Zi, Y. (2025). Trade From Space: Shipping Networks and the Global Implications of Local Shocks. *Review of Economics and Statistics*, pages 1–45.

Karimi-Mamaghan, M., Mohammadi, M., Pirayesh, A., Karimi-Mamaghan, A. M., and Irani, H. (2020). Hub-And-Spoke Network Design Under Congestion: A Learning Based Metaheuristic. *Transportation Research Part E: Logistics and Transportation Review*, 142:102069.

Kasahara, H. and Lapham, B. (2013). Productivity and the Decision to Import and Export: Theory and Evidence. *Journal of International Economics*, 89:297–316.

Khalil, M. and Weber, M.-D. (2022). Chinese Supply Chain Shocks. Discussion Paper 44/2022, Deutsche Bundesbank.

Kim, S.-J. and Shin, H. S. (2023). Theory of Supply Chains: A Working Capital Approach. Working Paper 1070, Bank for International Settlements.

Komaromi, A., Cerdeiro, D. A., and Liu, Y. (2022). Supply Chains and Port Congestion Around the World. Working Paper 2022/59, International Monetary Fund.

Li, Y., Bai, X., Wang, Q., and Ma, Z. (2022). A Big Data Approach to Cargo Type Prediction and Its Implications for Oil Trade Estimation. *Transportation Research Part E: Logistics and Transportation Review*, 165:102831.

Liu, E., Smirnyagin, V., and Tsyvinski, A. (2024). Supply Chain Disruptions and Supplier Capital in U.S. Firms. Working paper, SSRN.

Meng, Q., Wang, S., Andersson, H., and Thun, K. (2014). Containership Routing and Scheduling in Liner Shipping: Overview and Future Research Directions. *Transportation Science*, 48(2):265–280.

Merendino, A. and Monacelli, T. (2025). Supply Chain Uncertainty, Energy Prices, and Inflation. Mimeo. Bocconi University.

Michaillat, P. and Saez, E. (2015). Aggregate Demand, Idle Time, and Unemployment. *Quarterly Journal of Economics*, 130:507–569.

Michaillat, P. and Saez, E. (2022). An Economical Business-Cycle Model. *Oxford Economic Papers*, 74:382–411.

Mumtaz, H. and Zanetti, F. (2012). Neutral Technology Shocks and the Dynamics of Labor Input: Results From an Agnostic Identification. *International Economic Review*, 53:235–254.

Mumtaz, H. and Zanetti, F. (2015). Labor Market Dynamics: A Time-Varying Analysis. *Oxford Bulletin of Economics and Statistics*, 77:319–338.

Notteboom, T., Pallis, A., and Rodrigue, J.-P. (2022). *Port Economics, Management and Policy*. Routledge, 1st edition.

Rubio-Ramírez, J. F., Waggoner, D. F., and Zha, T. (2010). Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. *Review of Economic Studies*, 77:665–696.

Smets, F. and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97(3):586–606.

Smirnyagin, V. and Tsyvinski, A. (2022). Macroeconomic and Asset Pricing Effects of Supply Chain Disasters. Working Paper 30503, National Bureau of Economic Research.

Stopford, M. (2008). *Maritime Economics*. Routledge, 3rd edition.

Talley, W. K. (2009). *Port Economics*. Routledge, 1st edition.

Talley, W. K. and Ng, M. W. (2016). Port Multi-Service Congestion. *Transportation Research Part E: Logistics and Transportation Review*, 94:66–70.

Transportation Research Board Executive Committee (2006). *Critical Issues in Transportation*. The National Academies Press.

Uhlig, H. (2005). What Are the Effects of Monetary Policy on Output? Results From an Agnostic Identification Procedure. *Journal of Monetary Economics*, 52:381–419.

UNCTAD (2019). Review of Maritime Transport 2019. Technical report, United Nations.

Xu, L., Yu, Y., and Zanetti, F. (2025). The Adoption and Termination of Suppliers Over the Business Cycle. *Journal of Monetary Economics*, 151:103730.

Supplemental Appendices

The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory

Xiwen Bai[†], Jesús Fernández-Villaverde[‡], Yiliang Li[§], Francesco Zanetti[¶]

Contents

A Background on the Containerized Shipping Industry	A-3
B Testing Short-Run Rigidity in Containerized Shipping	A-5
C A Density-Based Spatial Clustering Algorithm	A-7
C.1 Methodology	A-7
C.2 Illustrative Cases	A-15
C.3 Weekly Indices of Port Congestion	A-22
D Discussion of Model Assumptions	A-24
E Proofs and Discussions Omitted from the Main Text	A-26
E.1 Proof of Proposition 1	A-26
E.2 Proof of Proposition 2	A-28
E.3 Slope and Curvature of Match Separation and Creation Schedules	A-31
E.4 Proof of Comparative Statics in Table 1	A-33
E.5 Convergence Dynamics	A-36
F External Data Sources for Baseline Estimation	A-42
G Robustness of Baseline Results	A-44
G.1 Prior Robustness	A-44
G.2 Price Levels vs. Inflation Rates in the SVAR	A-46
G.3 Invertibility Under an Expanded Information Set	A-51

[†] Bai: Tsinghua University, China. xiwenbai@mail.tsinghua.edu.cn. [‡] Fernández-Villaverde: University of Pennsylvania, U.S. jesusfv@econ.upenn.edu. [§] Li: University of International Business and Economics, China. yiliang_li@uibe.edu.cn. [¶] Zanetti: University of Oxford, U.K. francesco.zanetti@economics.ox.ac.uk.

G.4	Nonlinear Transmission by Shock Size	A-54
G.5	Different Lag Structures	A-57
G.6	Dropping the Linear Trend	A-61
G.7	Alternative Proxies for Activity and Prices	A-63
H	Identification Gains From Supply Chain Proxies	A-73
H.1	No Proxy	A-73
H.2	ACR Without Zero Restrictions	A-76
H.3	ACT	A-78
H.4	Trans-Pacific ACR	A-80
H.5	HARPEX	A-83
H.6	GSCPI	A-86
H.7	Supply Disruptions Index (SDI)	A-89
H.8	Goods Price Responses Across Proxies	A-92
References for Appendices		A-93

A. Background on the Containerized Shipping Industry

This appendix provides additional background on the containerized shipping industry, highlighting the institutional frictions and operating practices that shape its short-run dynamics.

Scheduling rigidities and port congestion. As noted in the main text, the shipping industry commonly follows a “hurry up and wait” practice in port calls (Du et al., 2015). For example, a vessel may depart the loading port at full speed to meet the originally requested time of arrival at the pilot boarding place (RTA PBP) on day 14. If severe congestion arises three days into the voyage and the RTA PBP is postponed to day 17, the vessel may not receive this update in time to adjust its speed. Even when such forewarnings are received, ships often maintain their speed because altering it could violate contractual obligations. As a result, vessels frequently “hurry” to arrive, only to “wait” at anchorage. This phenomenon is widespread in container shipping, where operational inefficiencies arise from mismatches between scheduled port calls and real-time port conditions.

Oil price, vessel speed, and port congestion. Fuel costs account for 50–60% of a liner shipping company’s operating expenses (Nottetboom, 2006), and vessel fuel consumption increases roughly with the cube of sailing speed (Li et al., 2016). As a result, container ship operators adjust sailing speeds in response to fluctuations in bunker oil prices. We estimate a strong negative relationship between oil prices and average vessel speed: a 1% increase in the Brent crude price reduces sailing speed by about 0.022%, a highly significant effect ($p = 0.004$). In contrast, when we regress congestion rates at the top 50 container ports worldwide on Brent crude prices (controlling for year and month fixed effects), we find no statistically significant relationship: the estimated coefficient is positive but small relative to its standard error, with a p -value of 0.313 and an R^2 of only 0.021.¹

This divergence reflects industry-wide institutional frictions. Higher oil prices encourage slower sailing to save fuel, but port congestion is shaped by scheduling rigidities and the “hurry up and wait” practice of port calls, which break the link between congestion dynamics and fuel prices. Thus, oil prices matter for cost-driven adjustments but not for congestion, which is governed by coordination failures and port inefficiencies.

¹Data on the Brent crude oil price are retrieved from the FRED database (mnemonic: POILBREUSDM). Data on average container ship speed, measured in knots, are retrieved from Clarksons. The statistically insignificant result for port congestion also holds when using the natural logarithm of average congestion time at the top 50 container ports.

Idle ships. According to the Clarksons Shipping Intelligence Network, idle ships are defined as vessels that (i) are not recorded with an average speed greater than one knot for seven days or more, (ii) are not identified as being under another status (such as laid-up, under repair, or in storage), and (iii) are either not subsequently recorded with an average speed greater than one knot for two or more consecutive days or have not moved more than 20 km.

Figure A.1 plots the share of idle container ships as a percentage of the global fleet. This proportion declined from about 6% to 3% at the end of 2017, hovered around 4–5% through 2018–2019, and then spiked to roughly 8.5% at the onset of the COVID-19 pandemic in early 2020. It subsequently fell sharply and stabilized at about 4%. The pandemic-driven surge in idle ships reflected active capacity management by shipping companies, which withheld capacity in response to the collapse in consumer and business demand (Li et al., 2025). To minimize the impact of this abrupt change on port-congestion estimates, we exclude idle ships when constructing the congestion indices (see Appendix C.1.1).

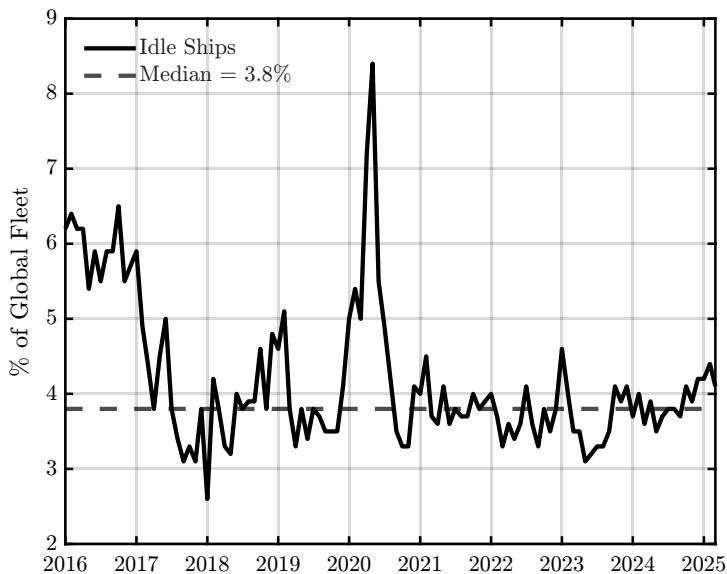


Figure A.1: Share of Idle Container Ships in the Global Fleet

B. Testing Short-Run Rigidity in Containerized Shipping

This appendix tests the short-run rigidity assumption in the containerized shipping industry. Specifically, we examine whether the ACR index responds on impact to a monetary policy shock using local projections (LPs). LPs offer a flexible model specification and are robust to the curse of dimensionality, making them well suited for capturing dynamic responses.

Our assessment involves estimating

$$\text{ACR}_{t+k} = \alpha_k + \beta_k \text{MP}_t + u_{k,t+k}, \quad (\text{B.1})$$

where $0 \leq k \leq K$ indexes the forecast horizon, ACR_{t+k} is the k -period-ahead value of the ACR index, MP_t is the monthly orthogonalized Bauer–Swanson monetary policy surprise published by the Federal Reserve Bank of San Francisco (Bauer and Swanson, 2023), α_k is a horizon-specific constant, and $u_{k,t+k}$ is the k -step-ahead forecast error.² The coefficients $\{\beta_k\}_{k=0}^K$ trace the response of ACR over K months to a one-standard-deviation policy shock, where the standard deviation is computed within our estimation sample.

We estimate Equation (B.1) via least squares (Jordà, 2005), using Newey–West standard errors (Newey and West, 1987) for β_k . As the forecast horizon k increases, the β_k estimates may become noisier, particularly in a small sample. To address this, we apply the smooth local projections (SLP) method (Barnichon and Brownlees, 2019), approximating β_k with a linear B-spline basis expansion:

$$\beta_k \approx \sum_{m=1}^M b_m B_m(k), \quad (\text{B.2})$$

where $B_m(k)$ are B-spline basis functions and b_m are parameters estimated using a penalized procedure that shrinks Equation (B.2) toward a quadratic polynomial in k .

We estimate the LPs beginning at horizon $k = 0$. Panel B.1a reports the impulse responses of the ACR index to a one-standard-deviation Bauer–Swanson monetary policy shock. The solid line denotes the median estimates, and the shaded areas show the 68% and 90% confidence bands. The on-impact response of the ACR index is close to zero and statistically insignificant, consistent

²We use the Bauer–Swanson series because it delivers precise estimates in recent samples. We rely on the monthly orthogonalized surprises from the “Monetary Policy Surprises” webpage maintained by the Federal Reserve Bank of San Francisco. As of our download, the monthly series runs through December 2023; accordingly, we estimate over January 2016–December 2023. This window still spans the key period of global supply chain disruptions in 2020–2022 and does not change our main conclusions.

with our assumption of short-run rigidity in the containerized shipping industry.

Beyond the immediate horizon, however, the ACR index exhibits a gradual decline over the medium run, with median responses turning more negative after six months and remaining below zero for up to two years. Although these medium-run effects are only marginally significant given the wide confidence bands, they suggest that monetary tightening may eventually dampen port congestion, likely through its effects on trade flows and shipping demand.

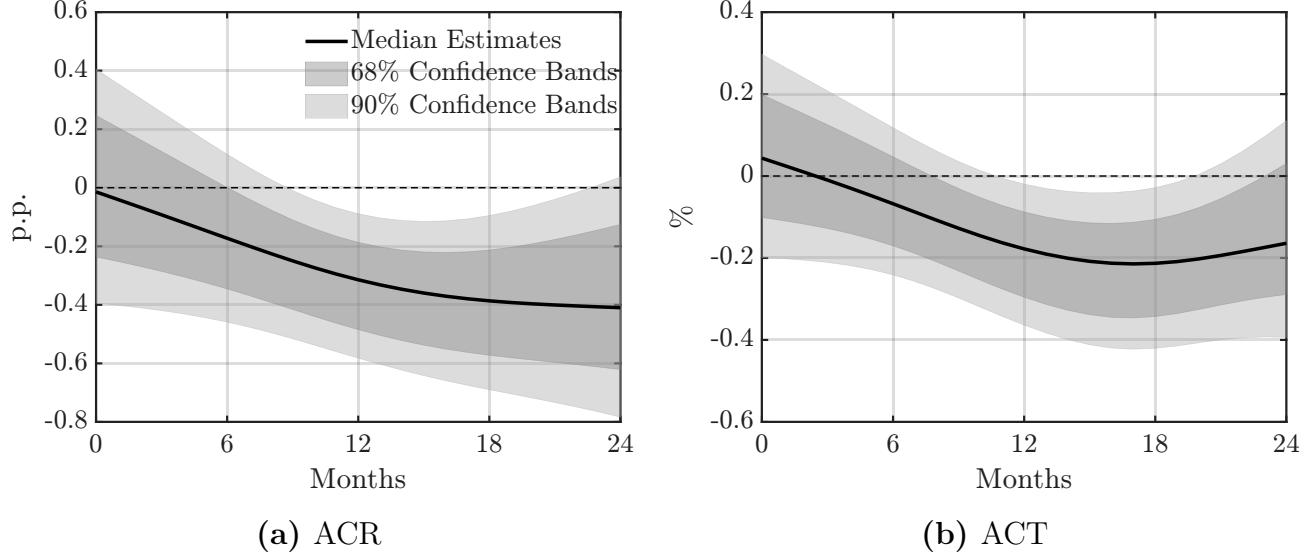


Figure B.1: Impulse Responses of Port Congestion Indices to a Monetary Policy Shock

Notes. The IRFs of the ACR and ACT indices to a one-standard-deviation Bauer–Swanson monetary policy shock, obtained from the LP model (B.1) estimated with the SLP method (Barnichon and Brownlees, 2019). Estimates are based on the sample period January 2016–December 2023.

Panel B.1b reports results using the Average Congestion Time (ACT) index, which serves as an alternative robustness check by measuring the average number of hours ships spend in anchorage before berthing. The ACT index exhibits a broadly similar pattern: no significant on-impact effect, followed by a modest but persistent decline over the medium run. This confirms that our findings are not an artifact of using a rate-based measure of congestion.

Taken together, these results support the assumption of short-run operational rigidity in the containerized shipping industry: neither the ACR nor the ACT index reacts contemporaneously to monetary policy shocks. At the same time, the medium-run declines observed in both indices are consistent with the view that monetary policy transmits to global shipping activity gradually, through its effects on demand for traded goods and shipping services.

C. A Density-Based Spatial Clustering Algorithm

This appendix provides details of our density-based spatial clustering algorithm —the iterative, multi-attribute, density-based spatial clustering of applications with noise (IMA-DBSCAN). The same details can be found in the companion paper (Bai et al., 2023).

We use this algorithm to estimate port congestion for the top 50 container ports worldwide.³ We first describe the methodology underlying the algorithm. We then present an illustrative case study for the Port of Ningbo-Zhoushan in China, showing its ability to identify both anchorage and berth locations —where other methods fall short. Finally, we construct weekly congestion indices —the ACR and ACT indices— to demonstrate that our measurement of port congestion is robust across time frequencies.

For completeness, the appendix also includes the pseudocode of IMA-DBSCAN to facilitate replication. The parameter values and ranges used in our estimation appear in Table C.1, and the AIS data inputs are publicly available. The algorithm can be adapted to compute congestion measures for ports beyond the top 50 container ports, and individual port-level congestion statistics are available upon request.

C.1. Methodology

As depicted in Figure C.1, the proposed IMA-DBSCAN algorithm has several distinctive features. Foremost among these is its two-tiered iterative structure. At the first level, we extract each container ship’s trajectory at each of the top 50 ports from the AIS data and apply traditional DBSCAN to filter out noise and cluster the ship’s mooring points (Ester et al., 1996). While this step identifies mooring areas, it does not sufficiently differentiate between the anchorage and berth areas of a port.

The second level addresses this limitation by applying a spatial-temporal DBSCAN (ST-DBSCAN; see Birant and Kut, 2007) to the resulting clusters. In this phase, we use an iterative procedure to determine a generalized and optimal parameter setting for the clustering algorithm. Another hallmark of IMA-DBSCAN is its integration of multiple attributes at the second level: beyond spatial data (e.g., coordinates), we incorporate non-spatial information (e.g., headings and timestamps) to enhance clustering accuracy. In what follows, we elaborate on the specifics of each level of IMA-DBSCAN.

³See <https://www.worldshipping.org/top-50-ports> (accessed June 15, 2022) for the full list of ports.

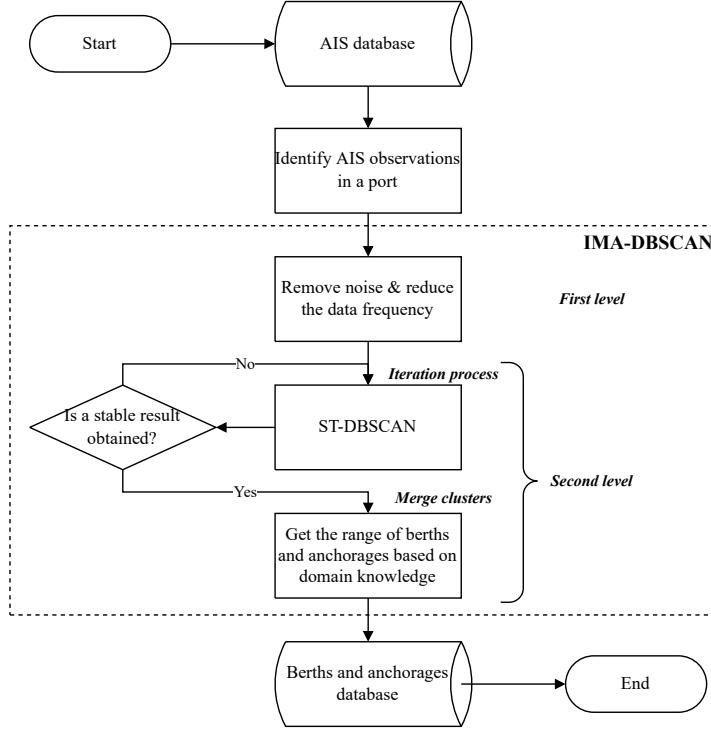


Figure C.1: Methodology Framework of IMA-DBSCAN

C.1.1. First Level: Data Preprocessing

While AIS data provide detailed information on ship positions, directly clustering these records to determine a port’s anchorage and berth areas is challenging. First, even when restricted to a specific port area and timeframe, the sheer volume of records means that running DBSCAN on the raw data leads to extended processing times. Second, the high incidence of incorrect AIS signal assignments can produce inaccurate clustering results, such as clusters that are not actual berths or that cover implausibly large areas. Third, if a ship remains in a port area for an extended period, the density of AIS records may cause DBSCAN to misclassify that stay as a cluster. Given these issues, preprocessing the AIS data is essential.

At the first level of IMA-DBSCAN, we filter the AIS records for each ship in the port area, focusing on observations with speeds below 1 knot. Such positions suggest that a ship is berthed, anchored, or in an unusual status (e.g., idle, laid-up, under repair, or in storage). We then tally these positions; if the count falls outside an acceptable range (e.g., fewer than 100 or more than 100,000), we classify the ship’s data as abnormal and exclude it from further analysis. Because a ship may dock at a port multiple times, we define a cutoff period Δt (e.g., 12 hours) to separate consecutive arrivals: if the gap between two stays exceeds Δt , we treat them as distinct port

calls. To streamline the dataset while maintaining consistency, we retain only the first data point for each hour.

For every port call, the refined positions are clustered using the traditional DBSCAN with parameters Eps and $MinPts$. We select an Eps value small enough to capture a ship’s mooring area and a $MinPts$ value sufficient to classify transient stops as noise. At this stage, AIS preprocessing is complete. The resulting samples are then passed to the second level of IMA-DBSCAN to identify a port’s anchorage and berth areas. For reference, the pseudocode for the first level of IMA-DBSCAN is provided in Algorithm 1.

C.1.2. Second Level: Multiple Attributes and Iteration

Information on headings. AIS data contain both spatial (geographical coordinates) and non-spatial (headings) information. Figure 1 in the main text illustrates the positions of a ship in a port alongside its headings. At a berth, headings are either aligned in the same direction or exactly opposite, whereas in an anchorage area they appear random with no discernible pattern. This observation matches real-world behavior: ships in anchorage often fail to maintain a consistent heading due to wind and wave variations.

Consequently, in the second level of IMA-DBSCAN, we incorporate heading information to improve estimation accuracy.⁴ Specifically, IMA-DBSCAN relies on three parameters: $Eps1$, $Eps2$, and $MinPts$. Here, $Eps1$ is the maximum geographical (spatial) distance between two points, $Eps2$ is the maximum non-spatial distance between two headings, and $MinPts$ is the minimum number of points within the thresholds defined by $Eps1$ and $Eps2$.

The geographical distance D is computed using the Haversine formula:

$$D[(x_1, x_2), (y_1, y_2)] = 2 \cdot R \cdot \arcsin \left[\sqrt{\sin^2 \left(\frac{x_1 - y_1}{2} \right) + \cos x_1 \cdot \cos y_1 \cdot \sin^2 \left(\frac{x_2 - y_2}{2} \right)} \right], \quad (C.1)$$

where the coordinates are in radians and $R = 6,371$ km is the Earth’s mean radius. The non-

⁴Non-spatial information is also useful for distinguishing between berths (see Algorithm 2). In our initial experiments, coordinates alone identified only the approximate locations of anchorage and berth areas, not the exact number of berths.

spatial distance Δh between two headings is defined as:

$$\Delta h(h_1, h_2) = \begin{cases} |h_1 - h_2|, & \text{if } |h_1 - h_2| \leq 180^\circ; \\ 360^\circ - |h_1 - h_2|, & \text{otherwise.} \end{cases} \quad (\text{C.2})$$

Using these distance measures, the neighbors of a point are those with a geographical distance less than $Eps1$ and a heading distance less than $Eps2$. A core point is defined as one with at least $MinPts$ neighbors, and clusters in IMA-DBSCAN consist only of these core points.

Iteration process. Since the geographical shapes of anchorage and berth areas vary widely across ports, and the boundaries of anchorage areas shift continuously with waves and winds, the parameters in IMA-DBSCAN must ideally vary across ports to achieve optimal results. To this end, we propose an iterative method for determining parameter values. Specifically, while $Eps2$ is fixed at 1° , $Eps1$ and $MinPts$ vary by port.

During the iteration, we define four intermediate variables: $Dist$, m , m' , and $NumC$. Here, $Dist$ is the average distance between a point in a cluster and the cluster's center; m is the number of points; m' is the number of noisy points (initialized to zero); and $NumC$ is the number of clusters.⁵ Using these variables, we update $Eps1$ and $MinPts$ as follows:

$$Eps1 = \alpha \cdot Dist, \quad MinPts = \beta \cdot \frac{m - m'}{NumC}.$$

Although α and β are not subject to explicit constraints, they must fall within a reasonable range to ensure both convergence and meaningful results. After evaluating performance under various parameter settings, we find $\alpha \in [0.4, 0.6]$ and $\beta \in [0.06, 0.1]$ to be appropriate. We also introduce $Dist_0$, which records the value of $Dist$ from the previous iteration (initialized to zero).

ST-DBSCAN is then executed iteratively. In each round, it operates with the current values of $Eps1$ and $MinPts$ and with $Eps2$ fixed at 1° . Its outputs classify points either into clusters or as noise. Based on these outputs, the intermediate variables, as well as $Eps1$ and $MinPts$, are updated, and the procedure is repeated. The process terminates when the difference $Dist - Dist_0$ is less than or equal to $\Delta Dist$ (e.g., 100 m).⁶

At convergence, each point is either assigned to a cluster or labeled as noise. We interpret

⁵At initialization, when no clusters exist, we treat all points as belonging to a single cluster. If all points are classified as noise, we set $NumC = 1$.

⁶The two-tiered structure and use of non-spatial information do not undermine the efficiency of IMA-DBSCAN. In practice, the values of $Eps1$ and $MinPts$ stabilize after only about five iterations, with further iterations yielding negligible changes.

the resulting cluster areas as berths and the areas with noisy points as anchorages.

Information on timestamps. After running ST-DBSCAN, we find that many clusters should be merged, as they often represent the same berth in reality. To improve the accuracy of berth identification, we merge clusters by using the timestamp information in the AIS data. Specifically, we first calculate the start and end times of each port call in each cluster. Since only one ship can dock at a berth at any given time, we then identify, for each cluster under consideration, its closest cluster and check for any overlap in docking times. If there is at least one overlap, the two clusters are treated as distinct berths. If there is no overlap, they are merged to represent a single berth (see Figure C.2 for an illustration).

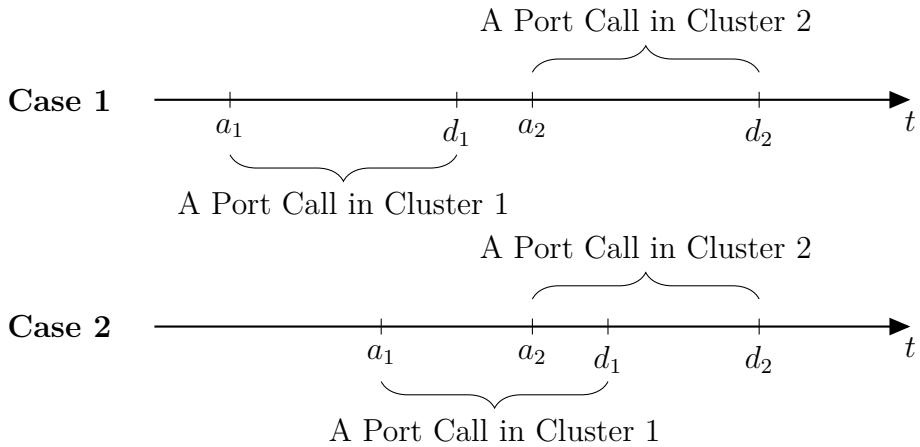


Figure C.2: Merging Clusters

Notes. Two scenarios that guide the criteria for merging clusters after executing ST-DBSCAN at the second level. Let a_1 and d_1 denote the arrival and departure times of a ship during a port call assigned to cluster 1. Similarly, let a_2 and d_2 represent the arrival and departure times for a port call assigned to cluster 2, the geographically closest cluster to cluster 1. In the first scenario, the docking intervals do not overlap, so clusters 1 and 2 are merged. In the second scenario, the docking intervals overlap, and the clusters are kept separate, since two ships cannot occupy the same berth simultaneously.

To distinguish between anchorage areas, we apply DBSCAN again to the points previously classified as noise. Here, the parameters Eps' and $MinPts'$ are chosen based on domain knowledge. We also remove clusters with fewer than N port calls, where N is set to align with official statistics (if any) on the minimum number of port calls recorded at the top 50 ports worldwide during the sample period. For reference, the pseudocode for the second level of IMA-DBSCAN is provided in Algorithms 2, 3, and 4.

Finally, the parameter values used in estimating port congestion are reported in Table C.1.

Table C.1: Parameter Values for IMA-DBSCAN

Parameter					
First Level		Δt	Eps	$MinPts$	
Value		12 hours	50 m	10	
Second Level	α	β	$\Delta Dist$	Eps'	$MinPts'$
Value	0.5	0.08	100 m	1,000 m	50
					N

Algorithm 1 Level 1 IMA-DBSCAN

Inputs:

$A_l = \{a_{1,l}, \dots, a_{n,l}\}$: the set of coordinates recorded in the AIS data for a ship l

$S_l = \{s_{1,l}, \dots, s_{n,l}\}$: the set of speeds recorded in the AIS data for a ship l

$T_l = \{t_{1,l}, \dots, t_{n,l}\}$: the set of timestamps recorded in the AIS data for a ship l

Outputs:

$D_l = \{d_{1,l}, \dots, d_{m,l}\}$: the coordinates of the first observation for each hour in B_l

$H_l = \{h_{1,l}, \dots, h_{m,l}\}$: the heading of the first observation for each hour in B_l

```

1: /* Data Preprocessing */                                 */
2:  $B_l = \{b_{1,l} \dots b_{k,l}\} \leftarrow$  the set of coordinates in  $A_l$  that indicate a speed less than 1 knot
3: /* Exception Identification */                         */
4: if  $|B_l| < 100$  or  $|B_l| > 100,000$  then
5:   Remove the data and stop                                ▷ The ship has an abnormal port call
6: else
7:   Continue
8: end if
9: /* DBSCAN Clustering */                                */
10:  $X \leftarrow b_{1,l}$ 
11: for  $i \leftarrow 2 : k$  do
12:   if  $t_i - t_{i-1} \leq \Delta t$  then
13:     Append  $b_{i,l}$  to  $X$ 
14:   else
15:     DBSCAN( $X, Eps, MinPts$ )
16:      $X \leftarrow \emptyset$ 
17:     Append  $b_{i,l}$  to  $X$ 
18:   end if
19: end for
20: Remove the observations labeled as noise from  $B_l$ 
21: Keep only the first observation for each hour in  $B_l$  ▷ Note that only  $m$  observations remain
     in  $B_l$  at this stage
22:  $D_l = \{d_{1,l}, \dots, d_{m,l}\} \leftarrow$  the coordinates of the first observation for each hour in  $B_l$ 
23:  $H_l = \{h_{1,l}, \dots, h_{m,l}\} \leftarrow$  the heading of the first observation for each hour in  $B_l$ 

```

Algorithm 2 Level 2 IMA-DBSCAN

Inputs:

$\mathbf{D} = \{D_1, \dots, D_L\}$: the set of coordinates for all ships after Level 1 IMA-DBSCAN

$\mathbf{H} = \{H_1, \dots, H_L\}$: the set of headings for all ships after Level 1 IMA-DBSCAN

$\mathbf{O} = \{\mathbf{D}, \mathbf{H}\} = \{o_1, \dots, o_M\}$: the combined set of coordinates and headings

Outputs:

C_{berth} : the set of clusters marked as berths

$C_{anchorage}$: the set of clusters marked as anchorages

```

1: /* Parameter Initialization */  

2:  $Dist \leftarrow$  the average distance between a point in  $\mathbf{D}$  and the center of the mass of  $\mathbf{D}$   

3:  $m \leftarrow |\mathbf{D}|$   

4:  $Eps1 \leftarrow \alpha \cdot Dist$   

5:  $MinPts \leftarrow \beta \cdot m$   

6: /* Iteration Process */  

7:  $Dist_0 \leftarrow 0$   

8: while  $Dist - Dist_0 > \Delta Dist$  km do  

9:   ST-DBSCAN( $\mathbf{O}, Eps1, Eps2 = 1^\circ, MinPts$ ) ▷ See function ST-DBSCAN  

10:   $Dist_0 \leftarrow Dist$   

11:   $Dist \leftarrow$  the average distance between a non-noisy point in  $\mathbf{D}$  and the center of the mass  

   of its assigned cluster  

12:   $m' \leftarrow |\text{noisy points in } \mathbf{O}|$   

13:   $NumC \leftarrow |\text{clusters in } \mathbf{O}|$   

14:   $Eps1 \leftarrow \alpha \cdot Dist$   

15:   $MinPts \leftarrow \beta \cdot (m - m') / NumC$   

16: end while  

17: /* Merging Clusters */  

18: Use the center of the mass of each cluster to calculate the distance in between  

19: for all clusters  $c$  in  $\mathbf{O}$  do  

20:    $c' \leftarrow$  the nearest cluster less than 500 m away from  $c$   

21:   if the docking times of  $c'$  and  $c$  do not overlap then  

22:     Replace the cluster label of  $c'$  with that of  $c$   

23:   end if  

24: end for  

25: /* Berth and Anchorage Detection */  

26:  $C_{berth} \leftarrow$  clusters in  $\mathbf{O}$   

27:  $C_{anchorage} \leftarrow$  DBSCAN(noisy points in  $\mathbf{O}, Eps', MinPts'$ )  

28: /* Exception Removal */  

29: for all clusters  $c$  in  $C_{berth}$  and  $C_{anchorage}$  do  

30:    $NumP \leftarrow$  the number of port calls in cluster  $c$   

31:   if  $NumP < N$  then  

32:     Remove  $c$   

33:   end if  

34: end for

```

Algorithm 3 ST-DBSCAN

Inputs:

$\mathcal{O} = \{o_1, \dots, o_M\}$: the combined set of coordinates and headings
 $Eps1$: maximum geographical (spatial) distance
 $Eps2$: maximum non-spatial distance
 $MinPts$: minimum number of points within the distance of $Eps1$ and $Eps2$

Outputs:

$C = \{c_1, \dots, c_M\}$: the set of clusters in \mathcal{O}

1: /* The codes are adapted from those in [Birant and Kut \(2007\)](#). */

2: **function** ST-DBSCAN($D, Eps1, Eps2, MinPts$)

3: $ClusterLabel = 0$

4: **for** $i \leftarrow 1 : m$ **do**

5: **if** o_i is not in a cluster **then**

6: $Y \leftarrow \text{RetrieveNeighbors}(o_i, Eps1, Eps2)$ *▷ See function RetrieveNeighbors*

7: **if** $|Y| < MinPts$ **then**

8: Mark o_i as noise

9: **else** *▷ Construct a new cluster*

10: $ClusterLabel \leftarrow ClusterLabel + 1$

11: **for** $j \leftarrow 1 : |Y|$ **do**

12: Mark all objects in Y with current $ClusterLabel$

13: **end for**

14: Push(all objects in Y)

15: **while** not IsEmpty() **do**

16: $CurrentObj = \text{Pop}()$

17: $Z \leftarrow \text{RetrieveNeighbors}(CurrentObj, Eps1, Eps2)$

18: **if** $|Z| \geq MinPts$ **then**

19: **for** all objects o in Z **do**

20: **if** o is not marked as noise **or** it is not in a cluster **then**

21: Mark o with current $ClusterLabel$

22: Push(o)

23: **end if**

24: **end for**

25: **end if**

26: **end while**

27: **end if**

28: **end if**

29: **end for**

30: $C = \{c_1, \dots, c_M\} \leftarrow$ the set of clusters in \mathcal{O}

31: **end function**

Algorithm 4 RetrieveNeighbors

Inputs:

o : an observation in O
 $Eps1$: maximum geographical (spatial) distance
 $Eps2$: maximum non-spatial distance

Outputs:

$Neighbors$: the set of neighbors for o

```
1: function RetrieveNeighbors( $o, Eps1, Eps2$ )
2:    $Neighbors \leftarrow \emptyset$ 
3:   for all observations  $o'$  in  $O$  do
4:      $Dist1 \leftarrow D(o, o')$                                  $\triangleright$  See Equation (C.1)
5:      $Dist2 \leftarrow \Delta h(o, o')$                        $\triangleright$  See Equation (C.2)
6:     if  $Dist1 \leq Eps1$  and  $Dist2 \leq Eps2$  then
7:       | Append  $o'$  to  $Neighbors$ 
8:     end if
9:   end for
10:  return  $Neighbors$ 
11: end function
```

C.2. Illustrative Cases

To demonstrate the capability of IMA-DBSCAN to accurately identify the anchorage and berth areas of a port —something other methods often fail to achieve— we apply the algorithm to the Port of Ningbo-Zhoushan in China, chosen for its intricate layout. Figure C.3a displays the first 50,000 AIS observations collected since January 1, 2020, within the port. The blue dots indicate the positions of low-speed container ships. Before applying IMA-DBSCAN, we mark approximate anchorage areas and berth locations using satellite images and nautical charts as benchmarks: red polygons denote anchorage areas, and yellow rectangles denote berth locations.⁷

Figure C.3b presents the clustering results of IMA-DBSCAN for Ningbo-Zhoushan, mirroring the layout in Figure C.3a for direct comparison. The clusters in Figure C.3b (colored red, yellow, blue, purple, cyan, and orange) closely match the anchorage areas in Figure C.3a.⁸

In Figure C.3e, we further highlight the locations of four terminals —Beilun, Daxie, Pukou, and Yuandong. Benchmarked against satellite maps, these identifications prove highly accurate:

⁷Our illustrative case of Ningbo-Zhoushan focuses on a one-month snapshot, as the first 50,000 AIS observations fall within January 2020. This identification remains representative in subsequent months because anchorage and berth areas do not change significantly in the short run. In practice, IMA-DBSCAN can be applied periodically to monitor potential changes in port layouts.

⁸For clarity, we also show the convex hulls of these clusters in Figure C.3d.

as shown in the top row of Figure C.4, each berth within the terminals is precisely delineated, with boundaries that align closely with reality.⁹

To benchmark IMA-DBSCAN against existing methods, we compare it with ST-DBSCAN.¹⁰ ST-DBSCAN is widely recognized in the literature as a leading spatial clustering algorithm capable of handling spatio-temporal data. Figure C.3c reports its results, which are less precise than those of IMA-DBSCAN. Although ST-DBSCAN generally identifies anchorage points, it misclassifies several high-density regions as berths, even when they are not. For instance, in the blue rectangle of Figure C.3f, ships that remained in an area for extended periods (possibly undergoing maintenance) are mistakenly treated as berths. Likewise, in the black rectangle, ST-DBSCAN labels points as berths that should instead be classified as mooring areas. As a result, ST-DBSCAN provides only a rough outline of anchorage areas and fails to identify berth locations with precision.

Figure C.4 offers a more granular comparison for the four Ningbo-Zhoushan terminals. ST-DBSCAN produces ambiguous results, with overlapping clusters that are spatially close but have divergent headings. Terminal boundaries can be discerned, yet individual berths are barely distinguishable. By contrast, IMA-DBSCAN delivers clusters that align cleanly with each berth. While tuning $MinPts$ or lowering $Eps1$ can improve ST-DBSCAN's accuracy, this requires continuous manual adjustment, which is difficult to apply consistently across ports. In contrast, the iterative design of IMA-DBSCAN automatically determines parameters, enabling accurate identification of both berths and anchorage areas.

Finally, we apply both algorithms to the Ports of Los Angeles and Long Beach (U.S.), Rotterdam (Netherlands), and Singapore. As shown in Figures C.5 through C.7, IMA-DBSCAN consistently outperforms ST-DBSCAN, delivering more accurate identification across all major container ports examined.

⁹Some blue dots in Figure C.3a do not correspond to any anchorage or berth in Figure C.3b, reflecting ships that anchored only briefly.

¹⁰For this comparison, the ST-DBSCAN parameters are set to $Eps1 = 2,500$ m, $Eps2 = 1^\circ$, and $MinPts = 100$, following [Ester et al. \(1996\)](#).

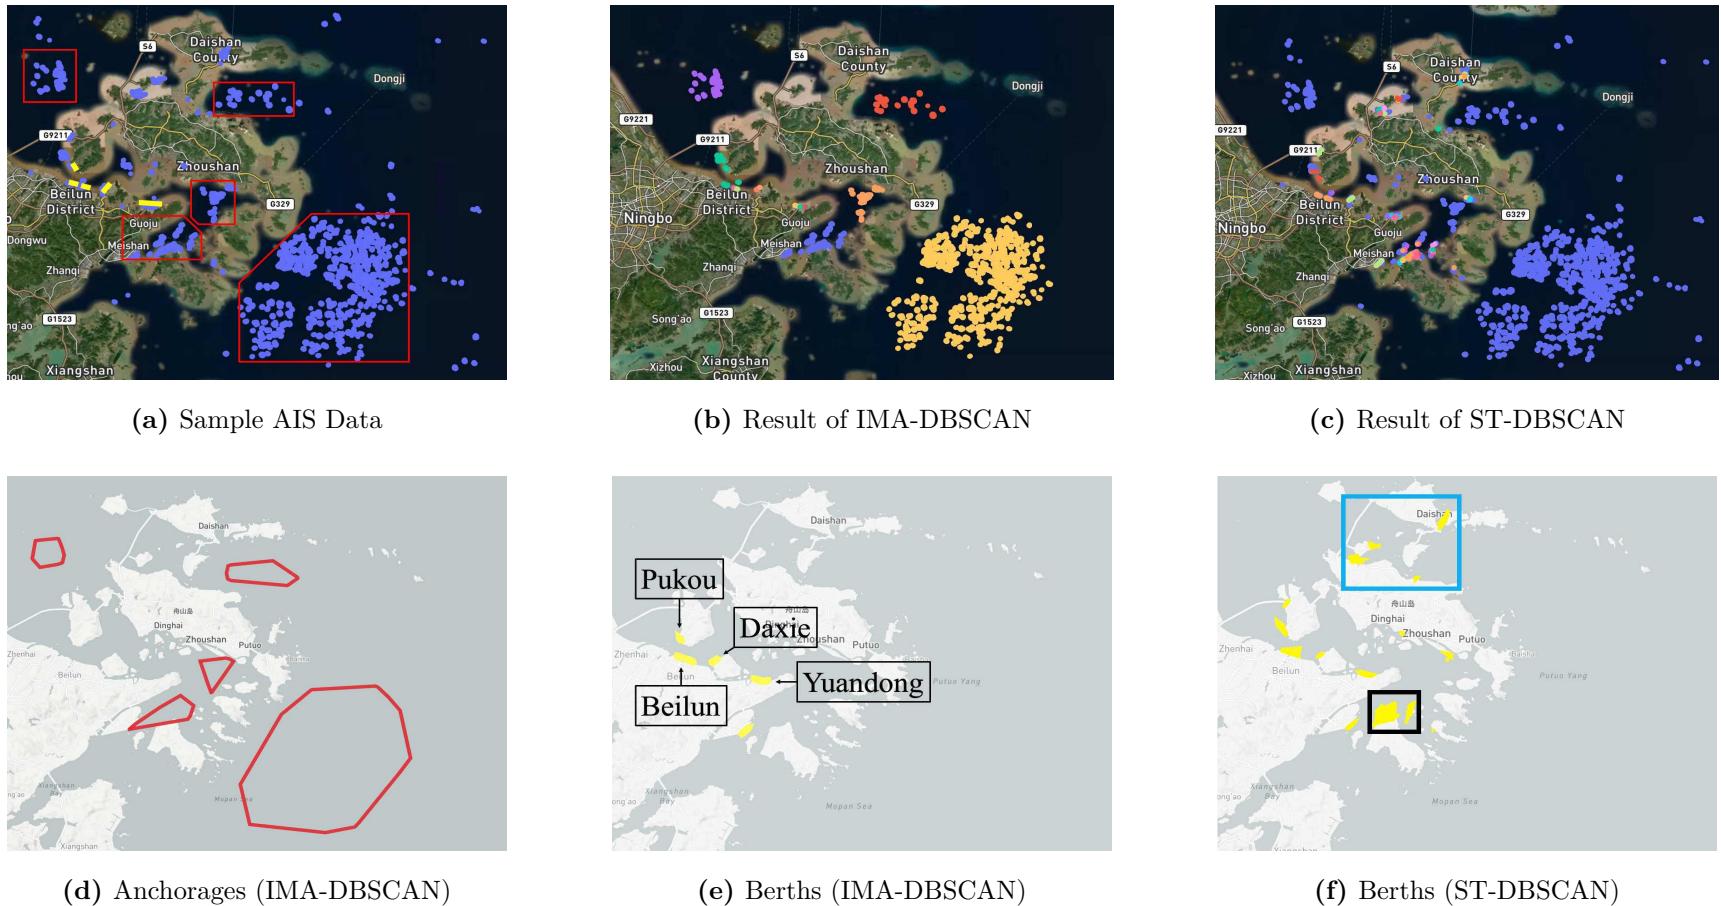
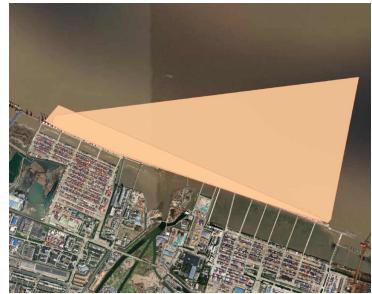


Figure C.3: Identification of Anchorage and Berth Areas in the Port of Ningbo-Zhoushan

Notes. In Panel (a), the sample data consist of the first 50,000 AIS observations recorded since January 1, 2020, within the Port of Ningbo-Zhoushan in China. These are shown as blue dots, covering coordinates from 121.60°E to 123.00°E and from 29.50°N to 30.35°N. As benchmarks, satellite maps and nautical charts are used to mark anchorage areas (red polygons) and berth locations (yellow rectangles). Panels (b) and (c) present the clustering results using IMA-DBSCAN and ST-DBSCAN, respectively. In Panel (b), the blue dots denote identified anchorage areas, whereas in Panel (c) they are classified as noise, outlining the general distribution of anchorages but without distinguishing individual ones. Panel (d) displays the anchorages from Panel (b) separately in red polygons, while Panel (e) shows the berths from Panel (b) separately in yellow, identifying the four terminals of Pukou, Daxie, Beilun, and Yuandong. Finally, Panel (f) depicts the approximate berth locations identified by ST-DBSCAN in yellow, with the blue and black rectangles indicating, respectively, noise misclassified as berths and anchorages mistaken for berths.



(a) Terminal Beilun



(b) Terminal Daxie



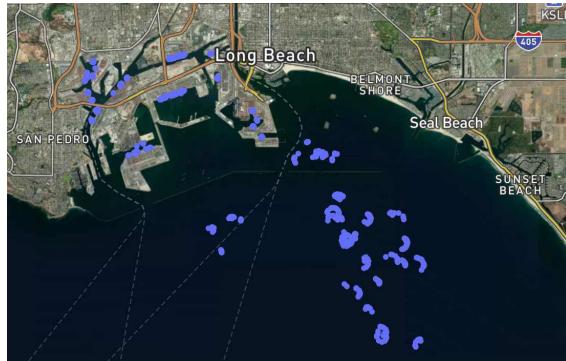
(c) Terminal Pukou



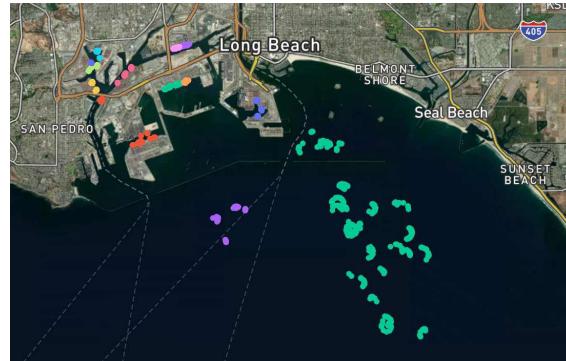
(d) Terminal Yuandong

Figure C.4: Detailed Results of Berth Identification: IMA-DBSCAN (Top Row) vs. ST-DBSCAN (Bottom Row)

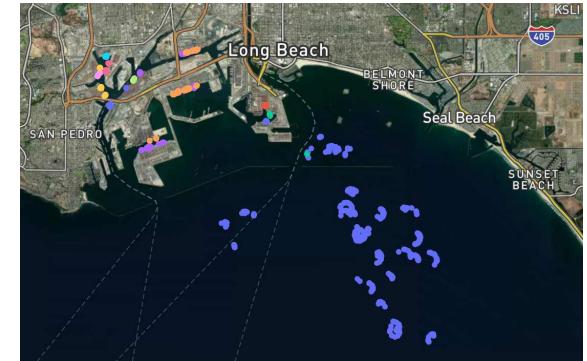
Notes. Detailed berth identification for the four terminals —Beilun, Daxie, Pukou, and Yuandong— within the Port of Ningbo-Zhoushan. Berths identified by IMA-DBSCAN appear in yellow (top row), and those identified by ST-DBSCAN appear in brown (bottom row).



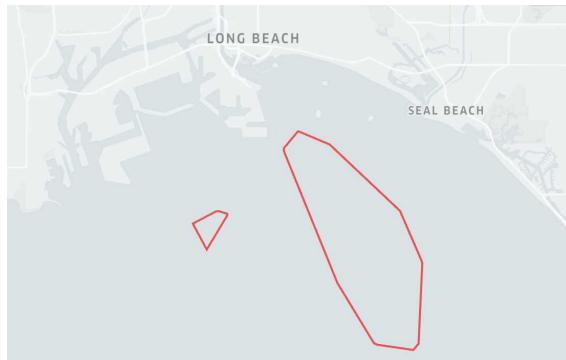
(a) Sample AIS Data



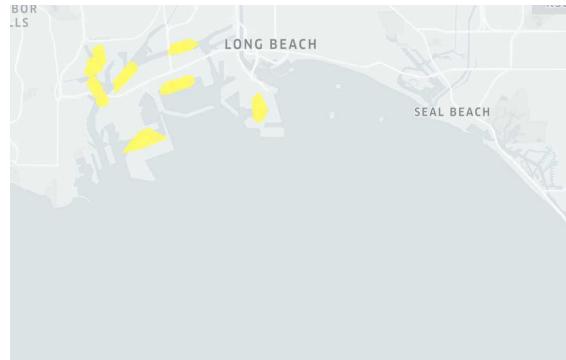
(b) Result of IMA-DBSCAN



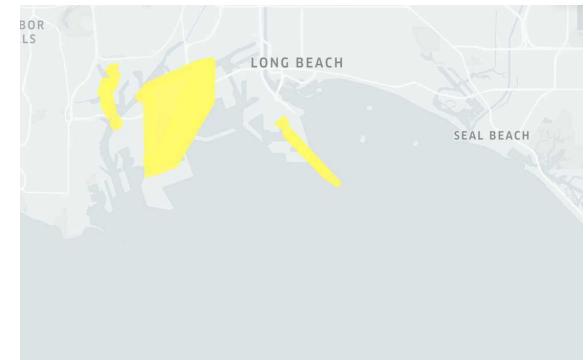
(c) Result of ST-DBSCAN



(d) Anchorages (IMA-DBSCAN)



(e) Berths (IMA-DBSCAN)



(f) Berths (ST-DBSCAN)

Figure C.5: Identification of Anchorage and Berth Areas in the Ports of Los Angeles and Long Beach

Notes. In Panel (a), the sample data consist of the first 50,000 AIS observations recorded since January 1, 2020, within the Ports of Los Angeles and Long Beach in the U.S., shown as blue dots. Panels (b) and (c) present the clustering results using IMA-DBSCAN and ST-DBSCAN, respectively. In Panel (b), the blue dots denote identified anchorage areas, whereas in Panel (c) they are classified as noise, outlining the general distribution of anchorages without distinguishing individual ones. Panel (d) displays the anchorages from Panel (b) separately in red polygons, and Panel (e) shows the berths from Panel (b) separately in yellow. Finally, Panel (f) depicts the approximate berth positions identified by ST-DBSCAN in yellow.



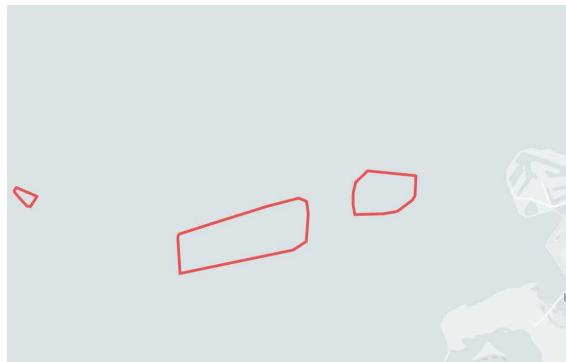
(a) Sample AIS Data



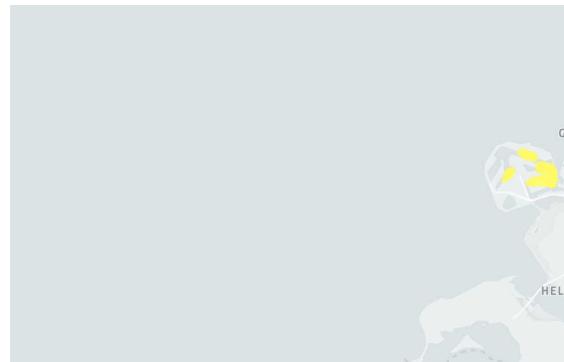
(b) Result of IMA-DBSCAN



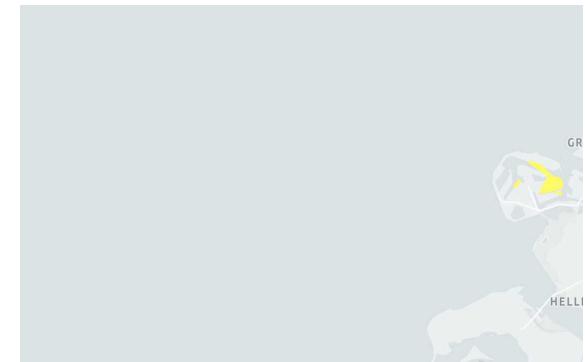
(c) Result of ST-DBSCAN



(d) Anchorages (IMA-DBSCAN)



(e) Berths (IMA-DBSCAN)



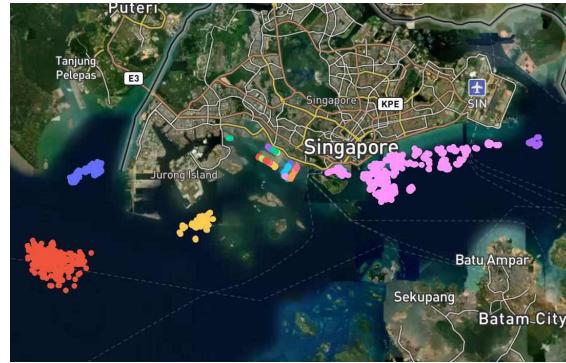
(f) Berths (ST-DBSCAN)

Figure C.6: Identification of Anchorage and Berth Areas in the Port of Rotterdam

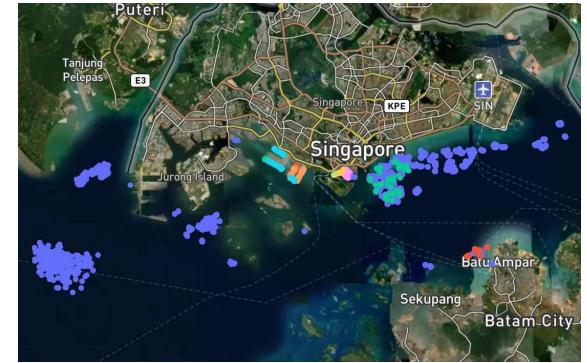
Notes. In Panel (a), the sample data consist of the first 50,000 AIS observations recorded since January 1, 2020, within the Port of Rotterdam in the Netherlands, shown as blue dots. Panels (b) and (c) present the clustering results using IMA-DBSCAN and ST-DBSCAN, respectively. In Panel (b), the blue dots denote identified anchorage areas, whereas in Panel (c) they are classified as noise, outlining the general distribution of anchorages without distinguishing individual ones. Panel (d) displays the anchorages from Panel (b) separately in red polygons, and Panel (e) shows the berths from Panel (b) separately in yellow. Finally, Panel (f) depicts the approximate berth positions identified by ST-DBSCAN in yellow.



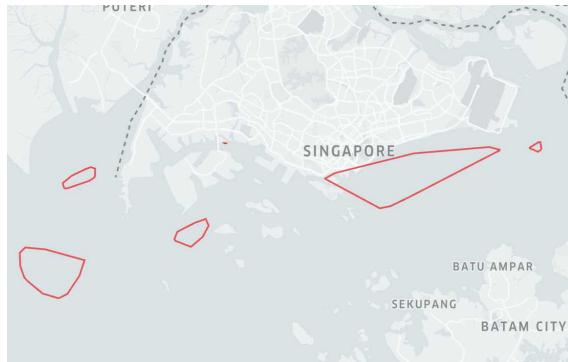
(a) Sample AIS Data



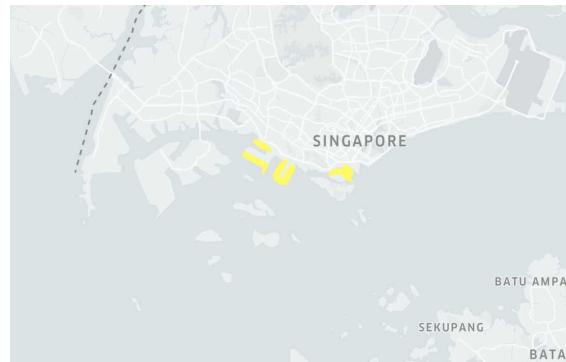
(b) Result of IMA-DBSCAN



(c) Result of ST-DBSCAN



(d) Anchorages (IMA-DBSCAN)



(e) Berths (IMA-DBSCAN)



(f) Berths (ST-DBSCAN)

Figure C.7: Identification of Anchorage and Berth Areas in the Port of Singapore

Notes. In Panel (a), the sample data consist of the first 50,000 AIS observations recorded since January 1, 2020, within the Port of Singapore, shown as blue dots. Panels (b) and (c) present the clustering results using IMA-DBSCAN and ST-DBSCAN, respectively. In Panel (b), the blue dots denote identified anchorages, whereas in Panel (c) they are classified as noise, outlining the general distribution of anchorages without distinguishing individual ones. Panel (d) displays the anchorages from Panel (b) separately in red polygons, and Panel (e) shows the berths from Panel (b) separately in yellow. Finally, Panel (f) depicts the approximate berth positions identified by ST-DBSCAN in yellow.

C.3. Weekly Indices of Port Congestion

The integration of high-frequency AIS data with our IMA-DBSCAN algorithm enables the construction of port congestion indices at frequencies higher than monthly updates. The AIS system processes over 2,000 reports per minute and can update information as often as every two seconds. Moreover, unlike traditional algorithms that require data sampled at fixed intervals, IMA-DBSCAN is uniquely flexible. Its streamlined design allows it to operate without predefined frequencies, making it especially well-suited to the highly variable and high-frequency nature of AIS data.

Figure C.8 shows the ACR index at a weekly frequency. Compared to the monthly ACR index in Figure 4 in the main text, the two series exhibit similar trends. However, the weekly series is substantially more volatile, since the number of ship visits to each port may vary considerably from week to week.¹¹

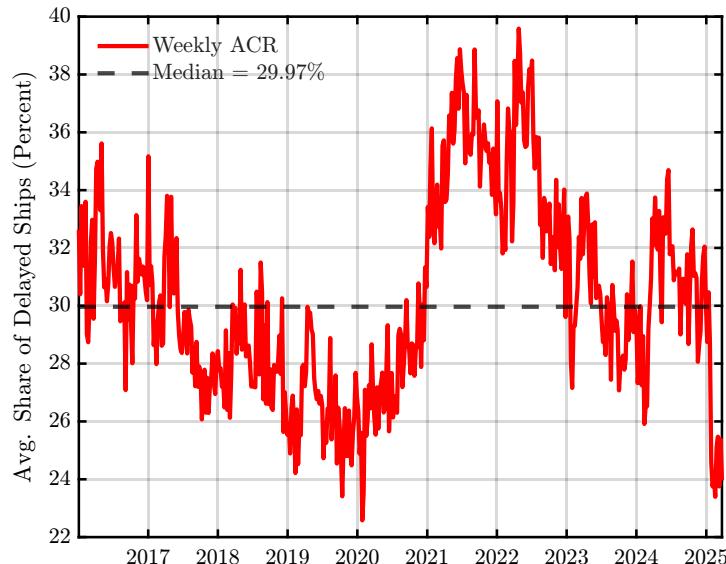


Figure C.8: Weekly ACR Index

In addition to the ACR index, we also introduce an alternative high-frequency congestion measure, the ACT index. Unlike the ACR index, which is rate-based, the ACT index measures the average number of hours a container ship spends waiting in a port's anchorage before docking at a berth, weighted by the relative number of ship visits to each of the top 50 container ports

¹¹The ACR index could also be constructed at daily or even hourly frequencies. However, because only a few ship visits occur at a given port each day or hour, the resulting series would be too volatile, with values of 0% and 100% occurring frequently.

worldwide:

$$\text{ACT}_t \equiv \sum_{p \in \mathcal{P}} \left[\frac{\text{Delayed}_{p,t} + \text{Undelayed}_{p,t}}{\sum_{p \in \mathcal{P}} (\text{Delayed}_{p,t} + \text{Undelayed}_{p,t})} \cdot \frac{\text{DelayHours}_{p,t}}{\text{Delayed}_{p,t} + \text{Undelayed}_{p,t}} \right],$$

where $\text{Delayed}_{p,t}$, $\text{Undelayed}_{p,t}$, and $\text{DelayHours}_{p,t}$ denote the number of delayed and undelayed ship visits, and the total hours ships spend in the anchorage areas of port p in week t , respectively.

Figure C.9 plots the weekly ACT index, which closely co-moves with the weekly ACR index over the sample period. In Appendix H.3, we also aggregate the ACT index to the monthly frequency and use it in our causality analysis. The results confirm that our identification strategy is robust to using ACT instead of ACR as the measure of global supply chain conditions.

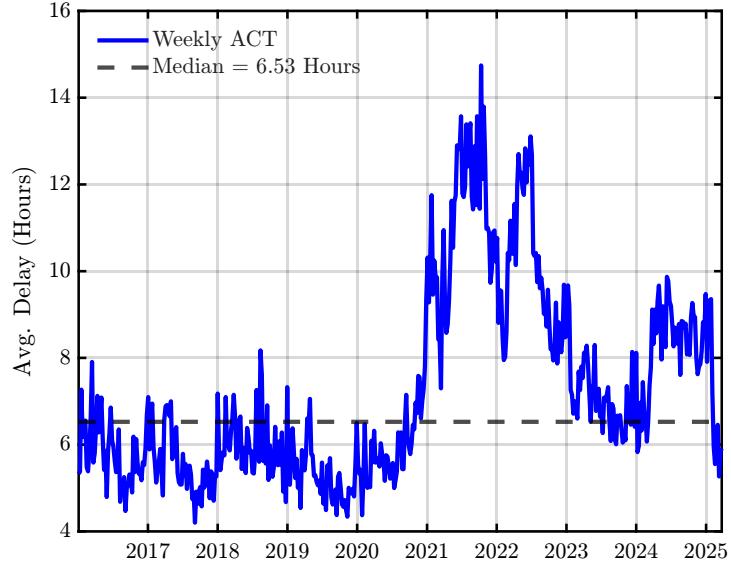


Figure C.9: Weekly ACT Index

Notes. The ACT index is a weighted average of the number of hours ships wait in a port's anchorage before docking at a berth, with the weights given by the relative number of ship visits to each port.

D. Discussion of Model Assumptions

This appendix discusses two critical assumptions in our model: search and matching frictions in the product market and endogenous separation of producer-retailer matches due to transportation costs. First, to represent the search and matching frictions in a tractable manner, we assume that a matching function governs the number of trades between producers and retailers. Second, to succinctly capture the decision-making process between a producer and a retailer when their trade is subject to a transportation cost, we assume that upon meeting, both parties endogenously separate once the idiosyncratic transportation cost exceeds a reservation threshold. We discuss each of these two assumptions in turn.

The matching function. There is ample literature on the sources of matching frictions in the product market, including the difficulty of locating and connecting with buyers across locations (Benguria, 2021; Krolkowski and McCallum, 2021; Lenoir et al., 2022), the costly acquisition of information about market conditions (Allen, 2014; Chaney, 2014), and informal trade barriers such as common language (Melitz and Toubal, 2014) and geography (Eaton and Kortum, 2002).

Common across all these theories is the presence of barriers to trade between producers and retailers, implying that not all unmatched producers engage in trade, and that not all retailers' sourcing visits succeed. We assume a simple constant-returns-to-scale matching function that summarizes how unmatched producers and retailers' visits are transformed into trades. This allows us to abstract from the complex matching process while preserving its key implication: unmatched producers trade only with probability $f(\theta)G(\bar{z}) < 1$, and retailers' visits are successful only with probability $q(\theta)G(\bar{z}) < 1$.

Endogenous separation on transportation cost. The separation margin in the product market can be modeled in the same way as in the labor market when workers face negative productivity shocks to their employment matches (Bils et al., 2011; Menzio and Shi, 2011; Fujita and Ramey, 2012). More concretely, producers face idiosyncratic transportation costs in their trading relationships with retailers, and sufficiently adverse draws lead to the termination of those relationships. This assumption is plausible only if there is convincing evidence that (i) transportation costs matter for the decision to trade, and (ii) there exists a threshold above which trading partners choose to sever their relationship.

The prediction that transportation cost affects the probability of trade has been extensively

corroborated by the trade literature. For example, evidence in [Rodrigue \(2020\)](#) shows that across all transport modes, a 10% increase in transportation costs reduces trade volumes by more than 20%. In maritime transportation specifically, [Brancaccio et al. \(2020\)](#) exploit tariff changes across the trade network to find that a 1% change in shipping costs generates roughly a 1% change in world trade value. Likewise, [Wong \(2022\)](#) estimates the elasticity of containerized trade with respect to freight rates using the round-trip effect as an instrument. For route i, j , she constructs a Bartik-style instrument to proxy for predicted trade on route j, i , and reports that a 1% increase in per-unit container freight rates decreases containerized trade value by 2.8% when dyad-by-product controls are included.

Both theory and observation point to the existence of a reservation transportation cost that trading partners consider when assessing whether a match is profitable. In practice, this threshold reflects the highest cost a firm can absorb when deciding to begin exporting or to maintain an existing trading relationship, covering expenses such as fuel, labor, and insurance. Evidence from plant-level data reinforces this idea. To account for the wide dispersion in export and import intensities across plants, [Kasahara and Lapham \(2013\)](#) extend the model of [Melitz \(2003\)](#) by allowing for heterogeneity in transportation costs. This modification delivers a natural self-selection mechanism: only plants with sufficiently low costs find it worthwhile to trade. These insights together support the notion of a reservation transportation cost below which trade becomes profitable.

A related threshold logic arises in the literature on transport infrastructure and exporting. For instance, [Naudé and Matthee \(2011\)](#) argue that firms require a minimum level of infrastructure in order to enter export markets, whereas further improvements have limited effects on export volumes. Because infrastructure affects transportation costs, this argument is consistent with the presence of a reservation transportation cost that firms weigh when making export decisions.

E. Proofs and Discussions Omitted from the Main Text

The following appendices include all long proofs that are omitted from the main text, as well as several important discussions on model dynamics.

E.1. Proof of Proposition 1

We first rewrite $\mathbb{E}_{z'} S(z')$. Using the definition of $S(z)$ in Equation (7), we derive the following expression for the total surplus:

$$S(z) = (p - z)l + (1 - \eta f(\theta)) \beta \mathbb{E}_{z'} S(z').$$

Subtracting Equation (11) from this expression yields $S(z) = (\bar{z} - z)l$. Substituting back into the expectation,

$$\begin{aligned} \mathbb{E}_{z'} S(z') &= l \int_0^{\bar{z}} (\bar{z} - z') dG(z') \\ &= l \left[(\bar{z} - z') G(z') \Big|_0^{\bar{z}} + \int_0^{\bar{z}} G(z') dz' \right] \\ &= l \int_0^{\bar{z}} G(z') dz', \end{aligned}$$

where the second line follows by integration by parts.

Replacing $\mathbb{E}_{z'} S(z')$ in Equations (11) and (12) gives the following rearranged match separation condition:

$$\mathbb{F}(p, \bar{z}, \theta) = 0 \quad \Rightarrow \quad p - \bar{z} + (1 - \eta f(\theta)) \beta \int_0^{\bar{z}} G(z') dz' = 0, \quad (\text{E.1})$$

and the rearranged match creation condition:

$$\mathbb{H}(\bar{z}, \theta) = 0 \quad \Rightarrow \quad \frac{\rho}{q(\theta)} - (1 - \eta) \beta \int_0^{\bar{z}} G(z') dz' = 0, \quad (\text{E.2})$$

where $G(\cdot)$ is the log-normal cumulative distribution function with density $g(\cdot)$.

Subsequently, we study the three-equation system given by the rearranged match separation condition (E.1), the match creation condition (E.2), and the retailer-household market clearing condition (18), i.e.,

$$\mathbb{L}(p, \bar{z}, \theta) = \frac{f(\theta)G(\bar{z})}{1 - G(\bar{z}) + f(\theta)G(\bar{z})} l - \chi^\varepsilon \frac{\mu}{p} = 0.$$

Assume \mathbb{F}, \mathbb{H} , and \mathbb{L} are C^1 and, at a solution $(p^*, \bar{z}^*, \theta^*)$, the partial derivatives satisfy:

$$\mathbb{F}_p = 1 > 0, \quad \mathbb{F}_{\bar{z}} = -1 + (1 - \eta f(\theta)) \beta G(\bar{z}) < 0, \quad \mathbb{F}_\theta = -\eta q(\theta)^{1+\xi} \beta \int_0^{\bar{z}} G(z') dz' < 0;$$

$$\mathbb{H}_{\bar{z}} = -(1 - \eta) \beta G(\bar{z}) < 0, \quad \mathbb{H}_\theta = \frac{\rho \theta^{\xi-1}}{q(\theta) (1 + \theta^\xi)} > 0;$$

$$\mathbb{L}_p = \frac{\chi^\varepsilon \mu}{p^2} > 0, \quad \mathbb{L}_{\bar{z}} = \frac{\phi [(\ln \bar{z} - \gamma) / \sigma] f(\theta) l}{\bar{z} \sigma (1 - G(\bar{z}) + f(\theta) G(\bar{z}))^2} > 0, \quad \mathbb{L}_\theta = \frac{(1 - G(\bar{z})) G(\bar{z}) q(\theta)^{1+\xi} l}{(1 - G(\bar{z}) + f(\theta) G(\bar{z}))^2} > 0,$$

and take $l > 0$, $\mu > 0$, $\chi > 0$, $p > 0$.

Consider the Jacobian with respect to (p, \bar{z}, θ) ,

$$J = \frac{\partial(\mathbb{F}, \mathbb{H}, \mathbb{L})}{\partial(p, \bar{z}, \theta)} = \begin{bmatrix} \mathbb{F}_p & \mathbb{F}_{\bar{z}} & \mathbb{F}_\theta \\ 0 & \mathbb{H}_{\bar{z}} & \mathbb{H}_\theta \\ \mathbb{L}_p & \mathbb{L}_{\bar{z}} & \mathbb{L}_\theta \end{bmatrix}.$$

At any solution $(p^*, \bar{z}^*, \theta^*)$,

$$\det J = \underbrace{\mathbb{F}_p}_{>0} \underbrace{(\mathbb{H}_{\bar{z}} \mathbb{L}_\theta - \mathbb{H}_\theta \mathbb{L}_{\bar{z}})}_{<0} + \underbrace{\mathbb{L}_p}_{>0} \underbrace{(\mathbb{F}_{\bar{z}} \mathbb{H}_\theta - \mathbb{F}_\theta \mathbb{H}_{\bar{z}})}_{<0} < 0,$$

because $\mathbb{H}_{\bar{z}} < 0 < \mathbb{H}_\theta$ and $\mathbb{L}_{\bar{z}}, \mathbb{L}_\theta > 0$ give the first bracket < 0 , while $\mathbb{F}_{\bar{z}} < 0$ and $\mathbb{F}_\theta < 0$ with $\mathbb{H}_{\bar{z}} < 0 < \mathbb{H}_\theta$ give the second bracket < 0 . Thus J is invertible and, by the Implicit Function Theorem, a solution is locally unique.

To obtain existence and global uniqueness, we reduce the system to a single equation in \bar{z} . Since $\mathbb{F}_p > 0$, for every (\bar{z}, θ) there is a unique $p = \lambda(\bar{z}, \theta)$ solving $\mathbb{F} = 0$, with:

$$\lambda_{\bar{z}} = -\frac{\mathbb{F}_{\bar{z}}}{\mathbb{F}_p} > 0, \quad \lambda_\theta = -\frac{\mathbb{F}_\theta}{\mathbb{F}_p} > 0.$$

Likewise, since $\mathbb{H}_{\bar{z}} < 0 < \mathbb{H}_\theta$, there is a unique $\theta = h(\bar{z})$ solving $\mathbb{H} = 0$ locally, with $h'(\bar{z}) = -\mathbb{H}_{\bar{z}}/\mathbb{H}_\theta > 0$. Define the scalar function:

$$\Lambda(\bar{z}) := \mathbb{L}(\lambda(\bar{z}, h(\bar{z})), \bar{z}, h(\bar{z})).$$

By the chain rule and the signs above,

$$\Lambda'(\bar{z}) = \underbrace{\mathbb{L}_p}_{>0} \left(\underbrace{\lambda_{\bar{z}}}_{>0} + \underbrace{\lambda_\theta}_{>0} \underbrace{h'(\bar{z})}_{>0} \right) + \underbrace{\mathbb{L}_{\bar{z}}}_{>0} + \underbrace{\mathbb{L}_\theta}_{>0} \underbrace{h'(\bar{z})}_{>0} > 0,$$

so Λ is strictly increasing on any interval where λ and h are defined.

Now, we bracket a root. As $\bar{z} \rightarrow 0^+$, $G(\bar{z}) \rightarrow 0$ and $\mathbb{L} \rightarrow -\chi^\varepsilon \mu/p < 0$, so $\Lambda(\bar{z}) < 0$ for small \bar{z} . As $\bar{z} \rightarrow +\infty$, $G(\bar{z}) \rightarrow 1$, so the first term of \mathbb{L} , i.e., $[f(\theta) G(\bar{z}) / (1 - G(\bar{z}) + f(\theta) G(\bar{z}))] l$, tends to l . Moreover,

$$\lambda_{\bar{z}} = -\mathbb{F}_{\bar{z}} = 1 - (1 - \eta f(\theta)) \beta G(\bar{z}) \geq 1 - \beta > 0,$$

because $G(\bar{z}) \leq 1$ and $1 - \eta f(\theta) \leq 1$. Hence $\lambda(\bar{z}, h(\bar{z}))$ grows at least linearly in \bar{z} , so $p = \lambda(\bar{z}, h(\bar{z})) \rightarrow +\infty$ and the second term $\chi^\varepsilon \mu/p \rightarrow 0$. Therefore $\Lambda(\bar{z}) \rightarrow l > 0$ as $\bar{z} \rightarrow +\infty$. By continuity and strict monotonicity of Λ , there exists a unique \bar{z}^* with $\Lambda(\bar{z}^*) = 0$. Setting $\theta^* = h(\bar{z}^*)$ and $p^* = \lambda(\bar{z}^*, \theta^*)$ then solves the original system.

Consequently, a steady state $(p^*, \bar{z}^*, \theta^*)$ exists and is globally unique over the relevant domain, and at that point the Jacobian is non-singular.

E.2. Proof of Proposition 2

Recall from Appendix E.1 that the rearranged match separation condition is:

$$\mathbb{F}(\bar{z}, \theta; p) : p = \bar{z} - (1 - \eta f(\theta)) \beta I(\bar{z}),$$

and the rearranged match creation condition is:

$$\mathbb{H}(\bar{z}, \theta) : \frac{\rho}{q(\theta)} = (1 - \eta) \beta I(\bar{z}),$$

where

$$I(\bar{z}) \equiv \int_0^{\bar{z}} G(z') dz',$$

and $G(\cdot)$ denotes the log-normal cumulative distribution function with density $g(\cdot)$.

Property 1. From \mathbb{H} , we obtain:

$$q(\bar{z}) = q(\theta(\bar{z})) = \frac{\rho}{(1 - \eta) \beta I(\bar{z})} \in (0, 1], \quad (\text{E.3})$$

which uniquely determines $\theta(\bar{z})$ because q is continuous and strictly decreasing in θ . Given that $f(\theta) = \theta q(\theta)$ and $f(\theta)^\xi + q(\theta)^\xi = 1$, this yields:

$$f(\bar{z}) = (1 - q(\bar{z})^\xi)^{\frac{1}{\xi}}, \quad (\text{E.4})$$

with $f(\bar{z}) \in [0, 1]$. Substituting into \mathbb{F} defines:

$$p(\bar{z}) = \bar{z} - (1 - \eta f(\bar{z})) \beta I(\bar{z}). \quad (\text{E.5})$$

Using $I_{\bar{z}}(\bar{z}) = G(\bar{z})$, $q_{\bar{z}}(\bar{z}) = -(q(\bar{z})/I(\bar{z})) G(\bar{z})$, and the identity $f(\theta)^\xi + q(\theta)^\xi = 1$, differentiating Equation (E.4) leads to:

$$f_{\bar{z}}(\bar{z}) = (1 - f(\bar{z})^\xi) f(\bar{z})^{1-\xi} \frac{G(\bar{z})}{I(\bar{z})}.$$

Subsequently, differentiating Equation (E.5) and performing a short algebraic rearrangement — rewriting $p(\bar{z}) = \bar{z} - \beta I(\bar{z}) + \eta \beta f(\bar{z}) I(\bar{z})$ and collecting terms in $f(\bar{z})$ and $f_{\bar{z}}(\bar{z})$ — yields:

$$\frac{dp}{d\bar{z}} = 1 - \beta G(\bar{z}) + \eta \beta G(\bar{z}) f(\bar{z})^{1-\xi}. \quad (\text{E.6})$$

Since $G(\bar{z}) \in (0, 1)$ and $f(\bar{z})^{1-\xi} \geq 0$, it follows that $dp/d\bar{z} > 1 - \beta > 0$. Hence $p(\bar{z})$ is strictly increasing, continuously differentiable, and therefore invertible. This establishes the existence, uniqueness, and C^1 regularity of the mapping $p \mapsto \bar{z}(p)$, and thus of $p \mapsto c_s^{ss}(p) = c_s^{ss}(\bar{z}(p))$. The minimal admissible \bar{z} is determined by $I(\bar{z}_{\min}) = \rho/[(1 - \eta)\beta]$, at which point $q(\bar{z}_{\min}) = 1$ and $f(\bar{z}_{\min}) = 0$.

Property 2. At \bar{z}_{\min} , Equation (E.5) yields $p_{\min} = \bar{z}_{\min} - \rho/(1 - \eta)$, while Equation (17) gives $c_s^{ss}(\bar{z}_{\min}) = 0$. As $\bar{z} \rightarrow +\infty$, we have $I(\bar{z}) \rightarrow +\infty$, implying $q(\bar{z}) \rightarrow 0$ and $f(\bar{z}) \rightarrow 1$ by Equations (E.3) and (E.4), and $G(\bar{z}) \rightarrow 1$. Consequently, $c_s^{ss}(\bar{z}) \rightarrow l$. From Equation (E.6), $dp/d\bar{z} \rightarrow 1 - \beta + \eta\beta > 0$, and thus $p \rightarrow +\infty$.

Property 3. Differentiating Equation (17) with respect to \bar{z} and defining $D(\bar{z}) \equiv 1 - G(\bar{z}) + f(\bar{z})G(\bar{z})$ yields:

$$\frac{dc_s^{ss}}{d\bar{z}} = \frac{f(\bar{z})g(\bar{z}) + G(\bar{z})(1 - G(\bar{z}))f_{\bar{z}}(\bar{z})}{D(\bar{z})^2} l.$$

Combining this expression with $dp/d\bar{z}$ from Equation (E.6) gives:

$$\frac{dc_s^{ss}}{dp} = \frac{f(\bar{z})g(\bar{z}) + G(\bar{z})(1 - G(\bar{z}))f_{\bar{z}}(\bar{z})}{D(\bar{z})^2(1 - \beta G(\bar{z}) + \eta \beta G(\bar{z}) f(\bar{z})^{1-\xi})} l > 0.$$

Moreover, as $\bar{z} \rightarrow +\infty$, we have $f(\bar{z}) \rightarrow 1$, $f_{\bar{z}}(\bar{z}) \rightarrow 0$, $D(\bar{z}) \rightarrow 1$, and $dp/d\bar{z} \rightarrow 1 - \beta + \eta\beta$. Hence,

$$\frac{dc_s^{ss}}{dp} \sim \frac{g(\bar{z})}{1 - \beta + \eta\beta} l \rightarrow 0, \quad (\text{E.7})$$

since $g(\bar{z}) \rightarrow 0$.

Property 4. Let $G_m \equiv G(\bar{z}_{\min})$ and $C \equiv \rho/[(1 - \eta)\beta]$. A first-order expansion of \mathbb{H} around

\bar{z}_{\min} yields:

$$I(\bar{z}) = C + G_m(\bar{z} - \bar{z}_{\min}) + o(\bar{z} - \bar{z}_{\min}),$$

and hence

$$q(\bar{z}) = \frac{C}{I(\bar{z})} = 1 - \frac{G_m}{C}(\bar{z} - \bar{z}_{\min}) + o(\bar{z} - \bar{z}_{\min}).$$

Since $f(\bar{z})^\xi = 1 - q(\bar{z})^\xi = \xi(G_m/C)(\bar{z} - \bar{z}_{\min}) + o(\bar{z} - \bar{z}_{\min})$, we obtain the local approximation:

$$f(\bar{z}) \sim \left(\frac{\xi G_m}{C} \right)^{\frac{1}{\xi}} (\bar{z} - \bar{z}_{\min})^{\frac{1}{\xi}}. \quad (\text{E.8})$$

Combining this with Equation (E.5) and noting that $I \approx C$ gives:

$$p - p_{\min} \sim (1 - \beta G_m)(\bar{z} - \bar{z}_{\min}) + \eta \beta C f(\bar{z}). \quad (\text{E.9})$$

If $\xi \in (0, 1)$, the first term in Equation (E.9) dominates. Using Equations (E.8) and (E.9) together with the first-order approximation of c_s^{ss} ,

$$c_s^{ss} \approx l \frac{G_m}{1 - G_m} f(\bar{z}),$$

we obtain:

$$c_s^{ss}(p) \sim \frac{l G_m (\xi G_m/C)^{\frac{1}{\xi}}}{(1 - G_m)(1 - \beta G_m)^{\frac{1}{\xi}}} (p - p_{\min})^{\frac{1}{\xi}},$$

which is locally convex.

If $\xi \geq 1$, the second term in Equation (E.9) dominates (or ties when $\xi = 1$), yielding a locally linear aggregate supply. In particular, if $\xi > 1$,

$$c_s^{ss}(p) \sim \frac{l G_m}{(1 - G_m) \eta \beta C} (p - p_{\min});$$

if $\xi = 1$,

$$c_s^{ss}(p) \sim \frac{l G_m^2}{(1 - G_m) C [1 - \beta G_m(1 - \eta)]} (p - p_{\min}).$$

Lastly, from Equation (E.7), as $\bar{z} \rightarrow +\infty$, we have:

$$\frac{dc_s^{ss}}{dp} \approx \kappa g(\bar{z}),$$

where $\kappa = l/(1 - \beta + \eta \beta) > 0$. Since the log-normal density g is unimodal and strictly decreasing in the upper tail, it follows that for sufficiently large \bar{z} (equivalently, large p), dc_s^{ss}/dp is strictly decreasing. Hence, $d^2 c_s^{ss}/dp^2 < 0$.

E.3. Slope and Curvature of Match Separation and Creation Schedules

Define:

$$I(\bar{z}) \equiv \int_0^{\bar{z}} G(z') dz',$$

where $G(\cdot)$ is the log-normal cumulative distribution function with strictly positive, continuous density $g(\cdot)$. Its derivatives are given by:

$$I_{\bar{z}}(\bar{z}) = G(\bar{z}) \in (0, 1), \quad I_{\bar{z}\bar{z}}(\bar{z}) = g(\bar{z}) > 0.$$

It is also useful to record the following first and second derivatives of $f(\theta)$ and $q(\theta)$ with respect to θ , which will be invoked repeatedly throughout the Appendix:

$$f_\theta(\theta) = \theta^{-\xi-1} (1 + \theta^{-\xi})^{-\frac{1}{\xi}-1} > 0, \quad f_{\theta\theta}(\theta) = -(1 + \xi) \theta^{-\xi-2} (1 + \theta^{-\xi})^{-\frac{1}{\xi}-2} < 0,$$

$$q_\theta(\theta) = -\theta^{\xi-1} (1 + \theta^\xi)^{-\frac{1}{\xi}-1} < 0, \quad q_{\theta\theta}(\theta) = (1 + \theta^\xi)^{-\frac{1}{\xi}-1} \theta^{\xi-2} \frac{1 - \xi + 2\theta^\xi}{1 + \theta^\xi}.$$

In particular, $q(\theta)$ is strictly convex for all $\theta > 0$ whenever $\xi \in (0, 1]$.

Match separation schedule. Using the rearranged match separation condition (E.1), by the Implicit Function Theorem, we have:

$$\frac{d\theta}{d\bar{z}} = -\frac{\mathbb{F}_{\bar{z}}}{\mathbb{F}_\theta} = \frac{\beta (1 - \eta f(\theta)) G(\bar{z}) - 1}{\beta \eta f_\theta(\theta) I(\bar{z})}.$$

Since $\beta \in (0, 1)$, $f(\theta) \in [0, 1)$, and $G(\bar{z}) < 1$,

$$\beta(1 - \eta f(\theta)) G(\bar{z}) < 1,$$

so the numerator of $d\theta/d\bar{z}$ is strictly negative while the denominator is strictly positive. Hence, $d\theta/d\bar{z} < 0$.

Differentiating once more gives:

$$\frac{d^2\theta}{d\bar{z}^2} = - \left[\mathbb{F}_{\bar{z}\bar{z}} + 2\mathbb{F}_{\bar{z}\theta} \frac{d\theta}{d\bar{z}} + \mathbb{F}_{\theta\theta} \left(\frac{d\theta}{d\bar{z}} \right)^2 \right] / \mathbb{F}_\theta,$$

where the partial derivatives are:

$$\mathbb{F}_{\bar{z}\bar{z}} = \beta(1 - \eta f(\theta)) g(\bar{z}) > 0, \quad \mathbb{F}_{\bar{z}\theta} = -\beta \eta f_\theta(\theta) G(\bar{z}) < 0,$$

$$\mathbb{F}_{\theta\theta} = -\beta \eta f_{\theta\theta}(\theta) I(\bar{z}) > 0, \quad \mathbb{F}_\theta = -\beta \eta f_\theta(\theta) I(\bar{z}) < 0.$$

Since $d\theta/d\bar{z} < 0$, each term inside the bracketed sum in the numerator is strictly positive, while the denominator is negative. Thus, $d^2\theta/d\bar{z}^2 > 0$, and the match separation schedule is strictly decreasing and strictly convex in (\bar{z}, θ) .

Match creation schedule. Recall from Appendix E.2 that the rearranged match creation condition is:

$$\mathbb{H}(\bar{z}, \theta) = \frac{\rho}{q(\theta)} - (1 - \eta)\beta I(\bar{z}) = 0.$$

From $q(\theta) = (1 + \theta^\xi)^{-1/\xi}$, we obtain the closed-form expression for $\theta(\bar{z})$:

$$\theta(\bar{z}) = \left\{ \left[\frac{(1 - \eta)\beta}{\rho} I(\bar{z}) \right]^\xi - 1 \right\}^{\frac{1}{\xi}},$$

which is well-defined on the interior domain:

$$I(\bar{z}) > \frac{\rho}{(1 - \eta)\beta},$$

or, equivalently, $\bar{z} \in [\bar{z}_{\min}, +\infty)$, where \bar{z}_{\min} is given by Equation (19) in the main text. Differentiating the closed form yields:

$$\frac{d\theta}{d\bar{z}} = \left[\frac{(1 - \eta)\beta}{\rho} \right]^\xi I(\bar{z})^{\xi-1} \left\{ \left[\frac{(1 - \eta)\beta}{\rho} \right]^\xi I(\bar{z})^\xi - 1 \right\}^{\frac{1}{\xi}-1} G(\bar{z}) > 0.$$

For curvature, write $\theta(\bar{z}) = h(I(\bar{z}))$ with:

$$h(I) = \left\{ \left[\frac{(1 - \eta)\beta}{\rho} \right]^\xi I^\xi - 1 \right\}^{\frac{1}{\xi}}.$$

Since $I_{\bar{z}}(\bar{z}) = G(\bar{z})$ and $I_{\bar{z}\bar{z}}(\bar{z}) = g(\bar{z})$, we obtain:

$$\frac{d^2\theta}{d\bar{z}^2} = \left[\frac{(1 - \eta)\beta}{\rho} \right]^\xi I(\bar{z})^{\xi-2} \theta(\bar{z})^{1-2\xi} [I(\bar{z}) g(\bar{z}) \theta(\bar{z})^\xi - (\xi - 1) G(\bar{z})^2].$$

The prefactor is strictly positive, so the sign of $d^2\theta/d\bar{z}^2$ is governed by the (second) bracketed term. If $\xi \in (0, 1]$, the bracketed term is strictly positive for all interior \bar{z} , implying $d^2\theta/d\bar{z}^2 > 0$.

If $\xi > 1$, the curvature is positive whenever

$$I(\bar{z}) g(\bar{z}) \theta(\bar{z})^\xi \geq (\xi - 1) G(\bar{z})^2.$$

Note that as $\bar{z} \rightarrow +\infty$, we have $g(\bar{z}) \rightarrow 0$ and $G(\bar{z}) \rightarrow 1$. In the far upper tail, this causes the left-hand side to eventually fall below the right-hand side, at which point the schedule becomes

locally concave. By contrast, on ranges where $g(\bar{z})$ remains sufficiently large and $G(\bar{z})$ is bounded away from 1, the inequality holds, implying convexity.

E.4. Proof of Comparative Statics in Table 1

We revisit the three-equation system defined in Appendix E.1. Let $x := (p, \bar{z}, \theta)^\top$ solve:

$$G(x; \mu, l, \gamma) = \begin{bmatrix} \mathbb{F}(p, \bar{z}, \theta; \gamma) \\ \mathbb{H}(\bar{z}, \theta; \gamma) \\ \mathbb{L}(p, \bar{z}, \theta; \mu, l, \gamma) \end{bmatrix} = 0,$$

and assume $\mathbb{F}, \mathbb{H}, \mathbb{L} \in C^1$.¹² Evaluating all partial derivatives at the steady state and adopting the sign pattern outlined in Appendix E.1,

$$\mathbb{F}_p > 0, \quad \mathbb{F}_{\bar{z}} < 0, \quad \mathbb{F}_\theta < 0; \quad \mathbb{H}_{\bar{z}} < 0, \quad \mathbb{H}_\theta > 0; \quad \mathbb{L}_p > 0, \quad \mathbb{L}_{\bar{z}} > 0, \quad \mathbb{L}_\theta > 0,$$

the Jacobian with respect to x is given by:

$$J := G_x = \begin{bmatrix} \mathbb{F}_p & \mathbb{F}_{\bar{z}} & \mathbb{F}_\theta \\ 0 & \mathbb{H}_{\bar{z}} & \mathbb{H}_\theta \\ \mathbb{L}_p & \mathbb{L}_{\bar{z}} & \mathbb{L}_\theta \end{bmatrix},$$

hence J is non-singular by the Implicit Function Theorem.

Comparative statics with respect to μ and l . Since only \mathbb{L} depends directly on μ and l ,

$$G_\mu = (0, 0, \mathbb{L}_\mu)^\top, \quad \mathbb{L}_\mu = -\frac{\chi^\varepsilon}{p} < 0;$$

$$G_l = (0, 0, \mathbb{L}_l)^\top, \quad \mathbb{L}_l = \frac{f(\theta)G(\bar{z})}{1 - G(\bar{z}) + f(\theta)G(\bar{z})} > 0,$$

by the Implicit Function Theorem,

$$x_\kappa = -J^{-1}G_\kappa, \quad \kappa \in \{\mu, l\},$$

and Cramer's rule yields:

$$p_\kappa = -\frac{(\mathbb{F}_{\bar{z}}\mathbb{H}_\theta - \mathbb{F}_\theta\mathbb{H}_{\bar{z}}) \mathbb{L}_\kappa}{\det J}, \quad \bar{z}_\kappa = \frac{(\mathbb{F}_p\mathbb{H}_\theta) \mathbb{L}_\kappa}{\det J}, \quad \theta_\kappa = -\frac{(\mathbb{F}_p\mathbb{H}_{\bar{z}}) \mathbb{L}_\kappa}{\det J}.$$

¹²Since the household's preference parameter for consuming goods, χ , enters the system in the same way as the money supply parameter, μ (through the aggregate demand function (16)), the comparative statics with respect to χ are identical to those for μ . We therefore omit the proof for brevity.

Given that:

$$\det J = \mathbb{F}_p (\mathbb{H}_{\bar{z}} \mathbb{L}_\theta - \mathbb{H}_\theta \mathbb{L}_{\bar{z}}) + \mathbb{L}_p (\mathbb{F}_{\bar{z}} \mathbb{H}_\theta - \mathbb{F}_\theta \mathbb{H}_{\bar{z}}) < 0, \quad \mathbb{F}_{\bar{z}} \mathbb{H}_\theta - \mathbb{F}_\theta \mathbb{H}_{\bar{z}} < 0, \quad \mathbb{F}_p \mathbb{H}_\theta > 0, \quad \mathbb{F}_p \mathbb{H}_{\bar{z}} < 0,$$

we obtain:

$$p_\mu > 0, \quad \bar{z}_\mu > 0, \quad \theta_\mu > 0; \quad p_l < 0, \quad \bar{z}_l < 0, \quad \theta_l < 0.$$

For the other variables of interest, it is straightforward to verify that their comparative statics with respect to μ and l are as follows:

$$c_\mu = \mathbb{L}_{\bar{z}} \bar{z}_\mu + \mathbb{L}_\theta \theta_\mu > 0, \quad r_\mu = \eta p_\mu + \eta \rho \theta_\mu > 0, \quad \frac{\partial (l - c)}{\partial \mu} = -c_\mu < 0;$$

$$c_l = -\mathbb{L}_p p_l > 0, \quad r_l = \eta p_l + \eta \rho \theta_l < 0, \quad \frac{\partial (l - c)}{\partial l} = \frac{1 - G(\bar{z})}{1 - G(\bar{z}) + f(\theta)G(\bar{z})} - \mathbb{L}_{\bar{z}} \bar{z}_l - \mathbb{L}_\theta \theta_l > 0.$$

Comparative statics with respect to γ , holding \bar{z} fixed. We focus on a short-run horizon in which producers' reservation transportation cost \bar{z} is fixed by long-term freight contracts, i.e., $\bar{z}_\gamma = 0$. Differentiating the three-equation system while treating \bar{z} as constant, we obtain:

$$\mathbb{F}_\gamma = -(1 - \eta f(\theta)) \beta \frac{1}{\sigma} \int_0^{\bar{z}} \phi\left(\frac{\ln z' - \gamma}{\sigma}\right) dz' < 0,$$

$$\mathbb{H}_\gamma = (1 - \eta) \beta \frac{1}{\sigma} \int_0^{\bar{z}} \phi\left(\frac{\ln z' - \gamma}{\sigma}\right) dz' > 0,$$

$$\mathbb{L}_\gamma = -\frac{f(\theta) l}{(1 - G(\bar{z}) + f(\theta)G(\bar{z}))^2} \frac{1}{\sigma} \phi\left(\frac{\ln \bar{z} - \gamma}{\sigma}\right) < 0,$$

and recall $\mathbb{H}_\theta > 0$, $\mathbb{L}_p > 0$, $\mathbb{L}_\theta > 0$. From $\mathbb{H}(\bar{z}, \theta; \gamma) = 0$ with $\bar{z}_\gamma = 0$:

$$\mathbb{H}_\theta \theta_\gamma + \mathbb{H}_\gamma = 0 \quad \Rightarrow \quad \theta_\gamma = -\frac{\mathbb{H}_\gamma}{\mathbb{H}_\theta} < 0.$$

From $\mathbb{L}(p, \bar{z}, \theta; \mu, l, \gamma) = 0$ with $\bar{z}_\gamma = 0$:

$$\mathbb{L}_p p_\gamma + \mathbb{L}_\theta \theta_\gamma + \mathbb{L}_\gamma = 0 \quad \Rightarrow \quad p_\gamma = -\frac{\mathbb{L}_\theta \theta_\gamma + \mathbb{L}_\gamma}{\mathbb{L}_p} > 0,$$

because $\mathbb{L}_\theta > 0$, $\theta_\gamma < 0$, $\mathbb{L}_\gamma < 0$, so the numerator $\mathbb{L}_\theta \theta_\gamma + \mathbb{L}_\gamma$ is negative, and $\mathbb{L}_p > 0$.

Equivalently, we can solve the system (\mathbb{H}, \mathbb{F}) for $(\theta_\gamma, p_\gamma)$. This yields the same $\theta_\gamma < 0$ and,

$$p_\gamma = \frac{-\mathbb{F}_\theta \theta_\gamma - \mathbb{F}_\gamma}{\mathbb{F}_p}.$$

To ensure consistency with the \mathbb{L} -based expression at the steady state, as $\mathbb{F}_p > 0$, a sufficient

condition for the \mathbb{F} -based numerator to be positive is:

$$-\mathbb{F}_\gamma \geq \mathbb{F}_\theta \theta_\gamma = -\frac{\mathbb{F}_\theta \mathbb{H}_\gamma}{\mathbb{H}_\theta}.$$

For the other variables of interest, the comparative statics of consumption c and spare capacity $l - c$ with respect to γ can be verified directly:

$$c_\gamma = -\mathbb{L}_p p_\gamma < 0, \quad \frac{\partial(l - c)}{\partial\gamma} = -c_\gamma > 0.$$

As for the wholesale price r , differentiating Equation (9) yields:

$$r_\gamma = \eta (p_\gamma + \rho \theta_\gamma).$$

Using the rearranged match separation condition (E.1),

$$p = \bar{z} - (1 - \eta f(\theta)) \beta I(\bar{z}), \quad I(\bar{z}) = \int_0^{\bar{z}} G(z') dz',$$

with $\bar{z}_\gamma = 0$ we get:

$$p_\gamma = -\beta [(1 - \eta f(\theta)) I_\gamma(\bar{z}) - \eta I(\bar{z}) f_\theta(\theta) \theta_\gamma], \quad I_\gamma(\bar{z}) := \frac{\partial I(\bar{z})}{\partial \gamma} < 0.$$

Next, from the rearranged match creation condition (E.2),

$$\frac{\rho}{q(\theta)} = (1 - \eta) \beta I(\bar{z}) \Rightarrow \rho = (1 - \eta) \beta I(\bar{z}) q(\theta),$$

and differentiating $\mathbb{H}(\bar{z}, \theta; \gamma) = 0$ with $\bar{z}_\gamma = 0$ gives:

$$\mathbb{H}_\theta \theta_\gamma + \mathbb{H}_\gamma = 0 \Rightarrow \theta_\gamma = -\frac{\mathbb{H}_\gamma}{\mathbb{H}_\theta} = \frac{(1 - \eta) \beta I_\gamma(\bar{z})}{\rho (-q_\theta(\theta)/q(\theta)^2)}.$$

Given that $f(\theta) = \theta q(\theta)$, $f_\theta(\theta) = q(\theta)^{1+\xi}$, and $q_\theta(\theta) = -\theta^{\xi-1} q(\theta)^{1+\xi}$, we have:

$$\theta_\gamma = \frac{(1 - \eta) \beta I_\gamma(\bar{z})}{\rho \theta^{\xi-1} q(\theta)^{\xi-1}} < 0.$$

Combine the pieces:

$$\begin{aligned} p_\gamma + \rho \theta_\gamma &= -\beta (1 - \eta f(\theta)) I_\gamma(\bar{z}) + \theta_\gamma [\beta \eta I(\bar{z}) f_\theta(\theta) + \rho] \\ &= -\beta (1 - \eta f(\theta)) I_\gamma(\bar{z}) + \theta_\gamma \beta I(\bar{z}) q(\theta) (1 - \eta + \eta q(\theta)^\xi) \\ &= \beta I_\gamma(\bar{z}) \left[- (1 - \eta f(\theta)) + \frac{1 - \eta + \eta q(\theta)^\xi}{\theta^{\xi-1} q(\theta)^{\xi-1}} \right], \end{aligned}$$

where in the second line we used $\rho = (1 - \eta)\beta I(\bar{z})q(\theta)$ and $f_\theta(\theta) = q(\theta)^{1+\xi}$, and in the third line we substituted the expression for θ_γ and canceled ρ . Using $f(\theta) = \theta q(\theta)$ and the identity $f(\theta)^\xi + q(\theta)^\xi = 1$, a short algebraic rearrangement of the bracketed term yields:

$$-(1 - \eta f(\theta)) + \frac{1 - \eta + \eta q(\theta)^\xi}{\theta^{\xi-1} q(\theta)^{\xi-1}} = \frac{1 - f(\theta)^{\xi-1}}{\theta^{\xi-1} q(\theta)^{\xi-1}}.$$

Therefore, the closed form of r_γ is given by:

$$r_\gamma = \eta(p_\gamma + \rho \theta_\gamma) = \eta\beta I_\gamma(\bar{z}) \frac{1 - f(\theta)^{\xi-1}}{\theta^{\xi-1} q(\theta)^{\xi-1}}.$$

With $I_\gamma(\bar{z}) < 0$ and $\theta^{\xi-1} q(\theta)^{\xi-1} > 0$ when $\theta \in (0, +\infty)$ (i.e., at an interior solution), we have:

$$\text{sign}(r_\gamma) = \begin{cases} + & \text{if } \xi \in (0, 1), \\ 0 & \text{if } \xi = 1, \\ - & \text{if } \xi > 1. \end{cases}$$

E.5. Convergence Dynamics

Next, we explore the convergence dynamics of our model from the initial steady state to the new steady state after unexpected and permanent shocks to aggregate demand, productive capacity, and the supply chain.

The dynamics of our model involve the evolution of three key endogenous variables —the number of matched producers ($x_{M,t}$), reservation transportation cost (\bar{z}_t), and product market tightness (θ_t)— following each exogenous shock.

Three equations govern these variables:

$$x_{M,t+1} = \left(1 + \theta_t^{-\xi}\right)^{-\frac{1}{\xi}} \Phi\left(\frac{\ln \bar{z}_t - \gamma}{\sigma}\right) + \left[1 - \left(1 + \theta_t^{-\xi}\right)^{-\frac{1}{\xi}}\right] \Phi\left(\frac{\ln \bar{z}_t - \gamma}{\sigma}\right) x_{M,t}, \quad (\text{E.10})$$

$$\frac{\chi^\varepsilon \mu}{x_{M,t} l} - \bar{z}_t + \left[1 - \eta \left(1 + \theta_t^{-\xi}\right)^{-\frac{1}{\xi}}\right] \beta \int_0^{\bar{z}_t} \Phi\left(\frac{\ln z' - \gamma}{\sigma}\right) dz' = 0, \quad (\text{E.11})$$

$$\rho \left(1 + \theta_t^\xi\right)^{\frac{1}{\xi}} - (1 - \eta) \beta \int_0^{\bar{z}_t} \Phi\left(\frac{\ln z' - \gamma}{\sigma}\right) dz' = 0, \quad (\text{E.12})$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. The other variables of interest, as listed in Table 1, are essentially functions of $x_{M,t}$, \bar{z}_t , and θ_t .

To analyze the convergence dynamics, we first compute the initial and new steady states of the

model before and after a change in the respective parameter. This requires setting $x_{M,t+1} = x_{M,t}$ and solving Equations (E.10), (E.11), and (E.12) simultaneously. We then assume the system starts at the initial steady state and examine the rate of convergence to the new steady state after each shock of interest. Lastly, we study whether the convergence is monotonic.

We calibrate the model at a weekly frequency using U.S. data, so that four periods in the theoretical model correspond roughly to one lag in the monthly SVAR. The baseline parameter values, which are standard in the literature, are reported in Table E.1. In particular, following Fernández-Villaverde et al. (2024), we set the producer's bargaining power, η , to 0.5, implying an equal split of the total surplus from matching. We also set $\sigma = 2$ to yield a steady-state spare capacity of 0.17. This value corresponds to the average spare capacity rate of 17% among the top five U.S. trading partners —Mexico, Canada, China, Germany, and Japan— during the pre-pandemic period 2016–2019, weighted by U.S. goods imports from each country in 2016. The import-weighted average spare capacity rate is estimated using Equation (23) in the main text, with data sources provided in Appendix F. Finally, the number of matched producers in the first iteration ($x_{M,1}$) is initialized at its steady-state value.

Table E.1: Baseline Calibration

Parameter	Description	Value
η	Bargaining power of producers	0.5
ρ	Fixed search cost	0.5
χ	Taste for consuming goods	1
ε	Elasticity of substitution between c and m/p	2
μ	Nominal money supply	10
l	Productive capacity	1
β	Discount factor	0.999
γ	Scale parameter of $G(\cdot)$	1
σ	Shape parameter of $G(\cdot)$	2
ξ	Elasticity of substitution between x_U and i_U	2

E.5.1. An Adverse Shock to Aggregate Demand

We first consider an adverse shock to aggregate demand, represented either by a 5% decrease in the money supply μ (from 10 to 9.5) or by a 5% decrease in the preference for consuming goods

χ (from 1 to 0.95). Tables E.2 and E.3 report the resulting convergence dynamics. Figures E.1 and E.2 illustrate the paths of consumption (equivalently, output) and price. In both scenarios, convergence to the new steady state is rapid, requiring only three iterations, because the model lacks a state variable, such as capital, that would generate persistence. The convergence dynamics are monotonic.

Table E.2: Convergence Dynamics: An Adverse Shock to Aggregate Demand ($\mu = 10 \rightarrow 9.5$)

Iteration	# Matched Producers	Consumption (or Output)	Price	Reservation Transportation Cost	Product Market	Wholesale Price	Spare Capacity
	x_M	c	p	\bar{z}	θ	r	$l - c$
	Initial Steady State	0.8281	0.8281	12.0753	18.1291	12.0247	18.1084
1	0.8206	0.8206	11.5773	17.1049	11.1775	17.1355	0.1794
2	0.8219	0.8219	11.5584	17.2839	11.3250	17.2524	0.1781
3	0.8217	0.8217	11.5617	17.2518	11.2986	17.2314	0.1783
New Steady State	0.8217	0.8217	11.5612	17.2567	11.3026	17.2346	0.1783

Notes. The values of the endogenous variables at the two steady states are obtained by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (E.10), (E.11), and (E.12) simultaneously. In this calculation, the parameter μ is set to 10 in the initial steady state and reduced to 9.5 in the new steady state (a 5% decrease). For the first iteration, $x_{M,1}$ is initialized at its steady-state value while μ is decreased to 9.5.

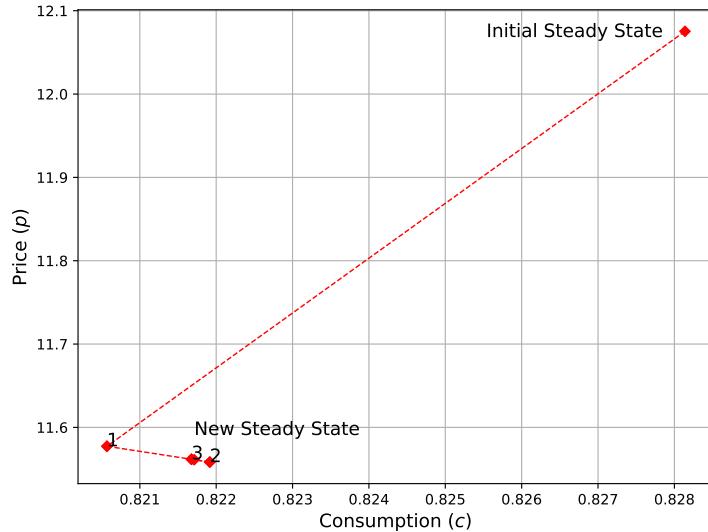


Figure E.1: Convergence Dynamics of Consumption and Price: An Adverse Shock to Aggregate Demand ($\mu = 10 \rightarrow 9.5$)

Notes. The iteration numbers are marked on the corresponding dots.

Table E.3: Convergence Dynamics: An Adverse Shock to Aggregate Demand ($\chi = 1 \rightarrow 0.95$)

Iteration	# Matched Producers	Consumption (or Output)	Price	Reservation Transportation Cost	Product Market Tightness	Wholesale Price	Spare Capacity
	x_M	c	p	\bar{z}	θ	r	$l - c$
	Initial Steady State	0.8281	0.8281	12.0753	18.1291	12.0247	18.1084
1	0.8128	0.8128	11.1038	16.1385	10.3850	16.2174	0.1872
2	0.8156	0.8156	11.0655	16.4846	10.6680	16.4420	0.1844
3	0.8151	0.8151	11.0726	16.4202	10.6153	16.4002	0.1849
New Steady State	0.8152	0.8152	11.0715	16.4301	10.6234	16.4067	0.1848

Notes. The values of the endogenous variables at the two steady states are obtained by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (E.10), (E.11), and (E.12) jointly. In this process, the parameter χ equals 1 in the initial steady state and is reduced to 0.95 (a 5% decrease) in the new steady state. For the first iteration, $x_{M,1}$ is set to its steady-state value, while χ is lowered to 0.95.

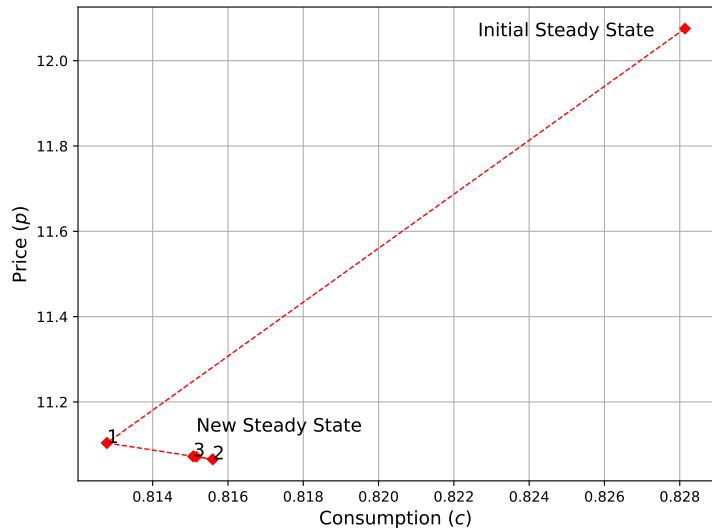


Figure E.2: Convergence Dynamics of Consumption and Price: An Adverse Shock to Aggregate Demand ($\chi = 1 \rightarrow 0.95$)

Notes. The iteration numbers are marked on the corresponding dots.

E.5.2. An Adverse Shock to Productive Capacity

Next, we examine a 5% reduction in the producers' fixed-factor endowment l from 1 to 0.95. As shown in Table E.4 and Figure E.3, convergence to the new steady state is rapid, requiring only three iterations, and the adjustment remains monotonic.

Table E.4: Convergence Dynamics: An Adverse Shock to Productive Capacity ($l = 1 \rightarrow 0.95$)

Iteration	# Matched Producers	Consumption (or Output)	Price	Reservation Transportation Cost	Product Market Tightness	Wholesale Price	Spare Capacity
	x_M	c	p	\bar{z}	θ	r	$l - c$
	Initial Steady State	0.8281	0.8281	12.0753	18.1291	12.0247	18.1084
1	0.8355	0.7937	12.5989	19.2144	12.9300	19.1392	0.1645
2	0.8343	0.7925	12.6176	19.0229	12.7697	19.0126	0.1657
3	0.8345	0.7927	12.6144	19.0547	12.7964	19.0337	0.1655
New Steady State	0.8344	0.7927	12.6149	19.0502	12.7925	19.0307	0.1656

Notes. The values of the endogenous variables at the two steady states are obtained by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (E.10), (E.11), and (E.12) simultaneously. In this process, the parameter l is 1 in the initial steady state and 0.95 in the new one (a 5% decrease). For the first iteration, $x_{M,1}$ is set to its initial steady-state value, and l is reduced to 0.95.

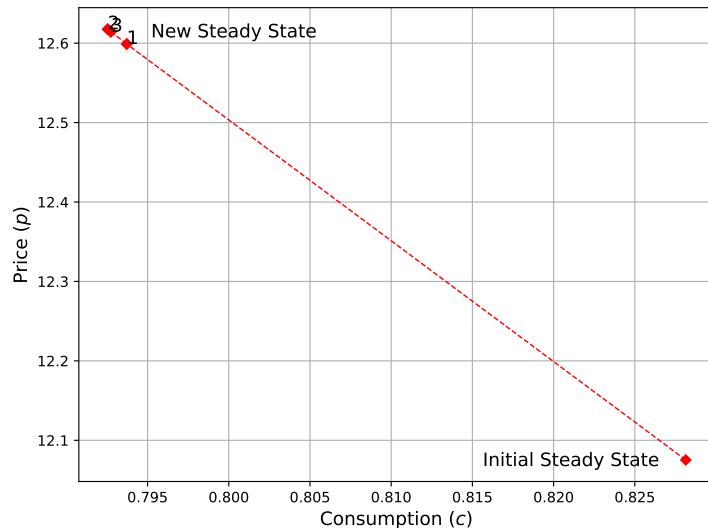


Figure E.3: Convergence Dynamics of Consumption and Price: An Adverse Shock to Productive Capacity ($l = 1 \rightarrow 0.95$)

Notes. The iteration numbers are marked on the corresponding dots.

E.5.3. An Adverse Shock to the Supply Chain

Lastly, we consider a 5% increase in the scale parameter of transportation costs, γ , from 1 to 1.05. As shown in Table E.5 and Figure E.4, the system converges to the new steady state within three iterations, and, as with the demand and capacity shocks, the adjustment is monotonic.

Table E.5: Convergence Dynamics: An Adverse Shock to the Supply Chain ($\gamma = 1 \rightarrow 1.05$)

Iteration	# Matched Producers	Consumption (or Output)	Price	Reservation Transportation Cost	Product Market Tightness	Wholesale Price	Spare Capacity
	x_M	c	p	\bar{z}	θ	r	$l - c$
	Initial Steady State	0.8281	0.8281	12.0753	18.1291	12.0247	18.1084
1	0.8208	0.8208	12.1834	18.0045	11.7737	18.0374	0.1792
2	0.8221	0.8221	12.1640	18.1875	11.9245	18.1569	0.1779
3	0.8219	0.8219	12.1675	18.1547	11.8974	18.1354	0.1781
New Steady State	0.8219	0.8219	12.1669	18.1596	11.9015	18.1387	0.1781

Notes. The values of the endogenous variables at the two steady states are obtained by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (E.10), (E.11), and (E.12) simultaneously. In this process, the parameter γ is set to 1 in the initial steady state and adjusted to 1.05 in the new steady state (i.e., a 5% increase). For the first iteration, $x_{M,1}$ is set to its initial steady-state value, and γ is increased to 1.05.

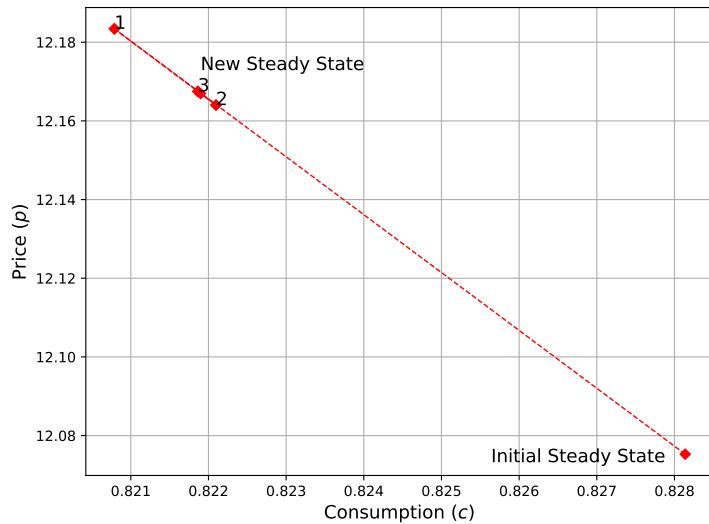


Figure E.4: Convergence Dynamics of Consumption and Price: An Adverse Shock to the Supply Chain ($\gamma = 1 \rightarrow 1.05$)

Notes. The iteration numbers are marked on the corresponding dots.

F. External Data Sources for Baseline Estimation

Table F.1 lists the external data sources used in the baseline SVAR estimation in Section 4, together with any adjustments applied. Below, we provide further details on the construction of spare capacity and product market tightness. The data used for robustness checks are described in subsequent appendices.

Spare capacity. In the main text, we define spare capacity as an import-weighted average of capacity utilization shortfalls across the top five U.S. trading partners (Mexico, Canada, China, Germany, and Japan). To implement this measure, we first construct monthly industrial production (IP) indices from official month-over-month growth rates, chaining forward from January 2016 and normalizing the starting value to 100. For Mexico and Japan, monthly capacity utilization rates are available and used as reported, except that Japan's index is rescaled so that its 2020 average equals 100 percent. For Canada, China, and Germany, where capacity utilization is reported only quarterly, we interpolate to monthly frequency using the Chow–Lin method ([Chow and Lin, 1971](#)) with the chained monthly IP series as indicators. We then apply fixed weights based on each country's share of U.S. imports in 2016 to obtain the aggregate spare capacity series used in estimation. We seasonally adjust the resulting series using the X-13ARIMA-SEATS algorithm provided by the U.S. Census Bureau.

Product market tightness. Tightness is defined as the ratio of U.S. manufacturers' new orders to the import-weighted value of spare capacity. To construct this variable empirically, we deflate the nominal series on U.S. manufacturers' new orders using the producer price index (PPI) for total manufacturing (manually seasonally adjusted), thereby expressing it in millions of constant 2005 U.S. dollars. The denominator, import-weighted spare capacity in U.S. dollar terms, is derived from country-level IP (converted into constant 2005 U.S. dollars using World Bank benchmarks) and capacity utilization rates, with the same interpolation procedures applied where necessary. These country-level series are aggregated using 2016 U.S. import shares as weights, and we seasonally adjust the final series. The resulting ratio provides the empirical measure of product market tightness used in the baseline SVAR estimation.

Table F.1: External Data Sources for Baseline Estimation

Variable	Mnemonic	Source	Notes on Construction/Adjustment
<i>U.S. Variables</i>			
Real PCE	PCEC96	U.S. Bureau of Economic Analysis	Raw series obtained directly from FRED.
PCE chain-type price index	PCEPI	U.S. Bureau of Economic Analysis	Raw series obtained directly from FRED.
Imports of goods: Mexico	IMPMX	U.S. Census Bureau	Customs basis, seasonally adjusted by authors (X-13ARIMA-SEATS).
Imports of goods: Canada	IMPCA	U.S. Census Bureau	Customs basis, seasonally adjusted by authors (X-13ARIMA-SEATS).
Imports of goods: China	IMPCH	U.S. Census Bureau	Customs basis, seasonally adjusted by authors (X-13ARIMA-SEATS).
Imports of goods: Germany	IMPGE	U.S. Census Bureau	Customs basis, seasonally adjusted by authors (X-13ARIMA-SEATS).
Imports of goods: Japan	IMPJP	U.S. Census Bureau	Customs basis, seasonally adjusted by authors (X-13ARIMA-SEATS).
Total manufacturers' new orders	AMTMNO	U.S. Census Bureau	Raw series obtained directly from FRED.
PPI: total manufacturing	PCUOMFGOMFG	U.S. Bureau of Labor Statistics	Seasonally adjusted by authors (X-13ARIMA-SEATS).
Import price index: all commodities	IR	U.S. Bureau of Labor Statistics	Seasonally adjusted by authors (X-13ARIMA-SEATS).
<i>Variables for Top U.S. Trading Partners</i>			
Capacity utilization rate: Mexico	N/A	National Institute of Statistics and Geography	Official monthly series..
Capacity utilization rate: Canada	N/A	Statistics Canada	Quarterly series interpolated to monthly frequency using IP.
Capacity utilization rate: China	N/A	National Bureau of Statistics of China	Quarterly series interpolated to monthly frequency using IP.
Capacity utilization rate: Germany	N/A	European Commission	Quarterly series interpolated to monthly frequency using IP.
Capacity utilization rate: Japan	N/A	Ministry of Economy, Trade and Industry	Official monthly series.
IP MoM growth: Mexico	N/A	National Institute of Statistics and Geography	Monthly series from official source.
IP MoM growth: Canada	N/A	Statistics Canada	Monthly series from official source.
IP MoM growth: China	N/A	National Bureau of Statistics of China	Monthly series from official source.
IP MoM growth: Germany	N/A	Federal Statistical Office	Monthly series from official source.
IP MoM growth: Japan	N/A	Ministry of Economy, Trade and Industry	Monthly series from official source.
IP: Mexico	N/A	World Bank	Annual series from World Development Indicators.
IP: Canada	N/A	World Bank	Annual series from World Development Indicators.
IP: China	N/A	World Bank	Annual series from World Development Indicators.
IP: Germany	N/A	World Bank	Annual series from World Development Indicators.
IP: Japan	N/A	World Bank	Annual series from World Development Indicators.

G. Robustness of Baseline Results

G.1. Prior Robustness

This appendix shows that the main conclusions from our baseline SVAR model remain robust when we use the prior-robust SVAR approach of [Giacomini and Kitagawa \(2021\)](#). This approach removes the need to specify a prior for the structural parameter based on the reduced-form parameter, the source of the asymptotic disagreement between Bayesian and frequentist inference. It does so by developing a class of priors that keeps a singular prior for the reduced-form parameter while allowing arbitrary conditional priors for the structural parameters, given the reduced-form parameter. This method strengthens the robustness of our SVAR analysis and ensures that our conclusions do not hinge on specific prior choices.

In practice, we apply their Algorithm 1 to approximate the set of posterior means and the associated robust credible regions for the IRFs of the selected endogenous variables in response to each structural shock. We make two modifications to Algorithm 1. First, in Step 2, to draw the orthonormal Q 's subject to Restrictions 1, 2, and 3, we use the QR decomposition method as in [Arias et al. \(2018\)](#) rather than the original linear projection approach. Both methods deliver comparable distributions of Q and similar computational costs. Second, we replace Step 3 of Algorithm 1 with Step 3' of Algorithm 2 to approximate the lower and upper bounds of the prior robust posterior means and the corresponding robust credible regions. These adjustments improve the precision and applicability of the procedure in our setting and yield more reliable approximations of posterior means and credible regions for the IRFs.

In Figures [G.1-G.3](#), the solid lines show the point-wise posterior medians, while the shaded areas display the 68% equal-tailed point-wise posterior probability bands. These bands are based on the baseline estimation sample from Section 4 in the main text. Dotted curves plot the set of prior-robust posterior means, and dashed-dotted curves indicate the associated 68% robust credible regions. All results are computed using 1,000 independent draws of the reduced-form parameters and 100,000 orthogonal matrix draws for each reduced-form parameter.

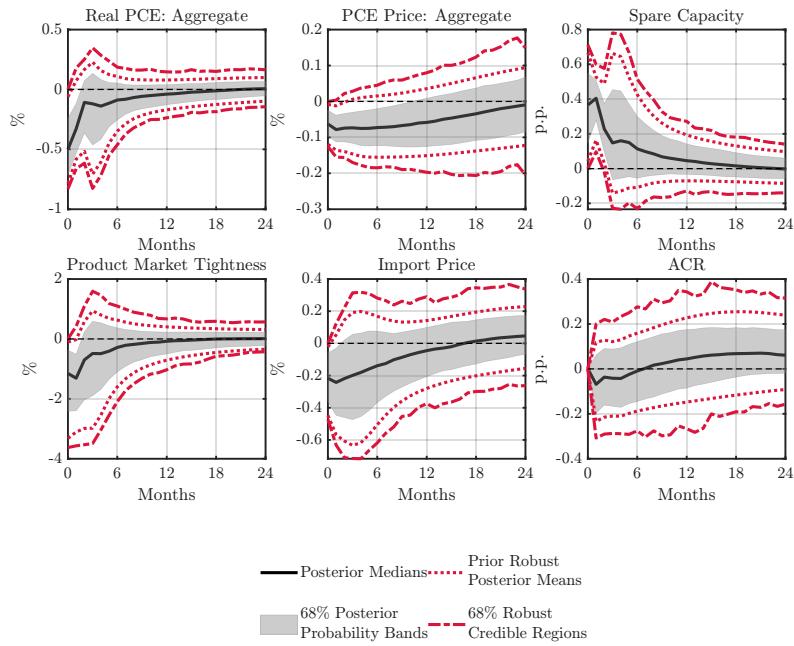


Figure G.1: IRFs to an Adverse Aggregate Demand Shock: Prior Robustness

Notes. IRFs to a one-standard-deviation adverse aggregate demand shock estimated using the prior-robust SVAR method proposed by [Giacomini and Kitagawa \(2021\)](#).

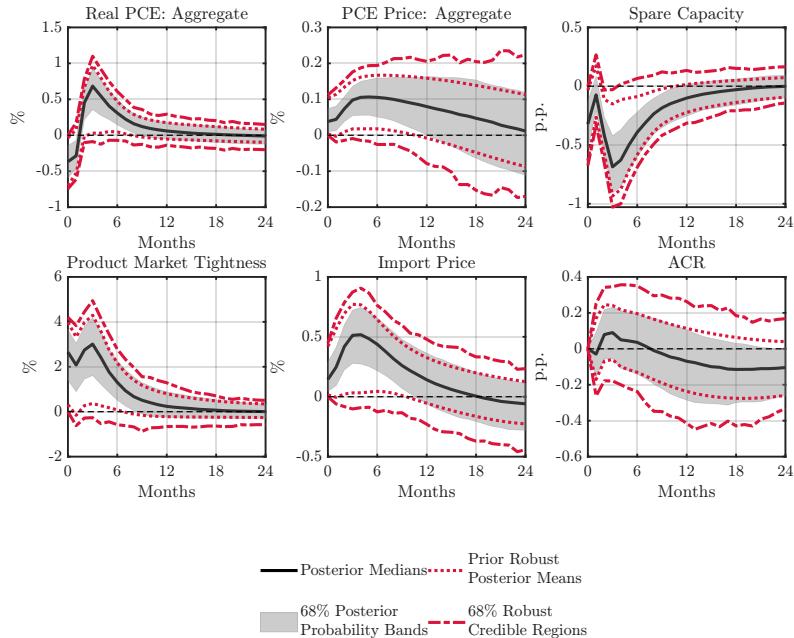


Figure G.2: IRFs to an Adverse Productive Capacity Shock: Prior Robustness

Notes. IRFs to a one-standard-deviation adverse productive capacity shock estimated using the prior-robust SVAR method proposed by [Giacomini and Kitagawa \(2021\)](#).

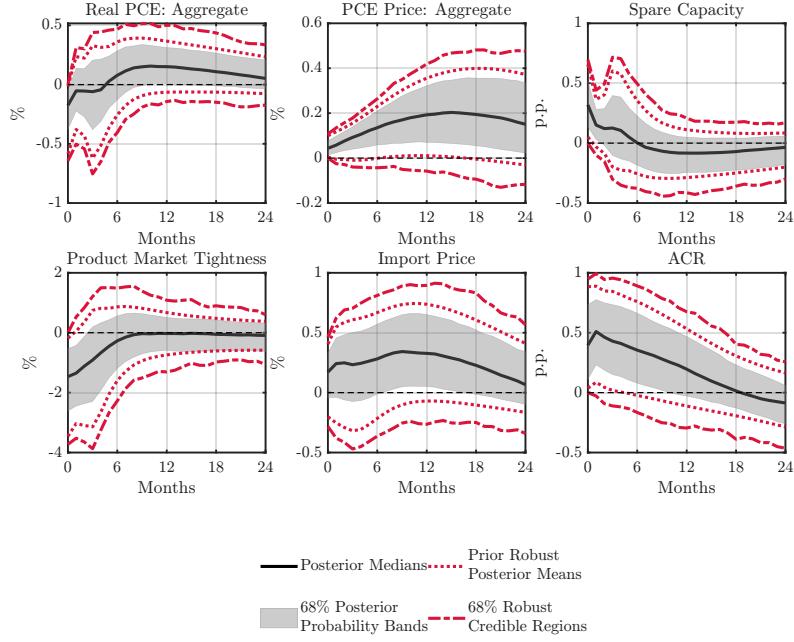


Figure G.3: IRFs to an Adverse Supply Chain Shock: Prior Robustness

Notes. IRFs to a one-standard-deviation adverse supply chain shock estimated using the prior-robust SVAR method proposed by [Giacomini and Kitagawa \(2021\)](#).

G.2. Price Levels vs. Inflation Rates in the SVAR

Next, we examine whether our results depend on using price levels rather than inflation rates.¹³ Starting from the baseline SVAR, we replace the PCE chain-type price index and the import price index with their monthly inflation rates using log differences, leaving the variables, priors, identification (sign and zero restrictions), and lag length unchanged. Let π_t denote inflation. For any horizon k , the implied response of the log price level equals the cumulative sum of the inflation IRFs, $\text{IRF}_{\Delta \ln p}(k) = \sum_{j=0}^k \text{IRF}_\pi(j)$, so the inflation and level specifications are directly comparable.

Figures G.4-G.6 illustrate this comparison. For both PCE and import prices, the cumulated inflation responses (orange) closely track the corresponding price-level IRFs from the levels SVAR (gray) across horizons, with similar impact, hump shape, and reversion. Figures G.7-G.9 further show that real-side responses (real PCE, spare capacity, product market tightness, and ACR) remain qualitatively unchanged relative to the baseline: same signs, similar peak timing, and comparable persistence. Thus, moving from price levels to inflation does not alter our results.

¹³We use the levels specification in the baseline because our theoretical framework in Section 3 delivers sign predictions for the price level (a one-off increase rather than a permanent change in inflation).

Next, we compare FEVDs across the two specifications. Figure G.10 reports those from the inflation-based SVAR. Relative to the price-level specification in Figure 10, for PCE inflation the supply chain disturbance accounts for the largest share of the forecast error variance at medium to long horizons in both models, with demand and capacity shocks playing smaller roles. For import price inflation, all three shocks contribute roughly equally, with demand shocks marginally more important. For real PCE, spare capacity, and product market tightness, either demand or capacity shocks dominate at short to medium horizons in both specifications, while the supply chain contribution remains moderate but rises gradually with the horizon. For the ACR, the supply chain shock is the primary driver throughout. Overall, the ordering and magnitudes of contributions are robust to whether the model is estimated in levels or in inflation rates.

Finally, the inflation responses are consistent with canonical DSGE evidence. For instance, [Smets and Wouters \(2007\)](#) show that a positive demand shock raises output and inflation simultaneously (Figure 2), while a positive productivity shock raises output and temporarily lowers inflation (Figure 7). Our adverse demand and capacity shocks generate the mirror image, and the short-run magnitude of the PCE inflation response in our data is comparable to theirs.

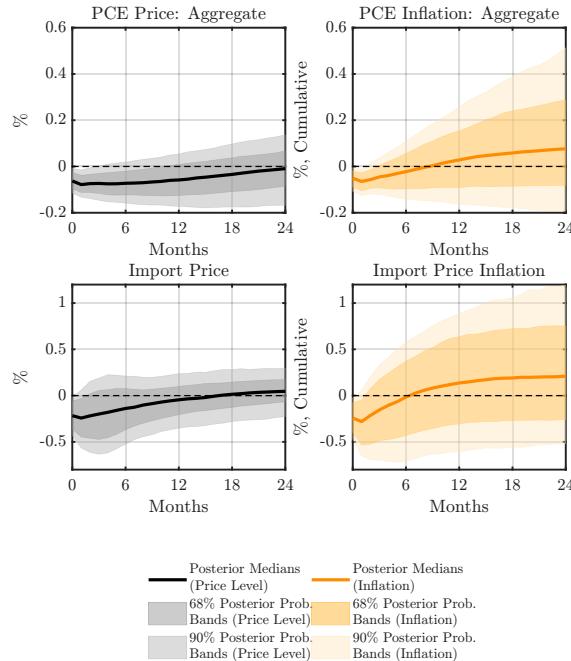


Figure G.4: Price-Level vs. Cumulated-Inflation IRFs: Adverse Aggregate Demand Shock

Notes. Price-level (gray) and cumulated inflation (orange) IRFs for PCE and import prices after an adverse aggregate demand shock.

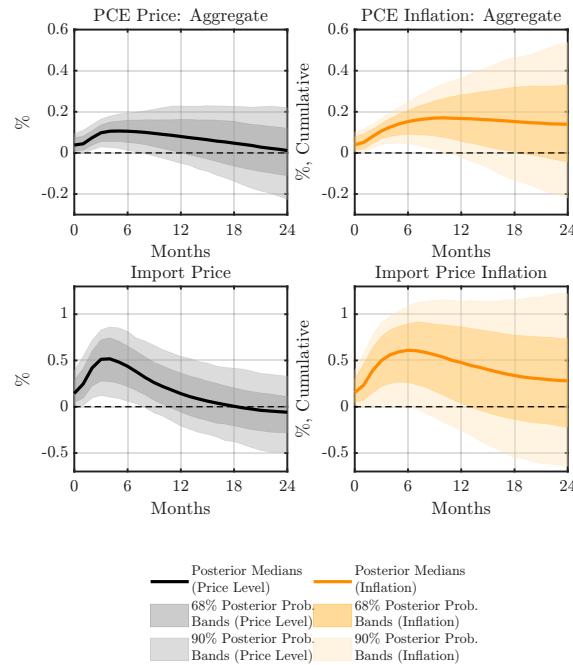


Figure G.5: Price-Level vs. Cumulated-Inflation IRFs: Adverse Productive Capacity Shock

Notes. Price-level (gray) and cumulated inflation (orange) IRFs for PCE and import prices after an adverse productive capacity shock.

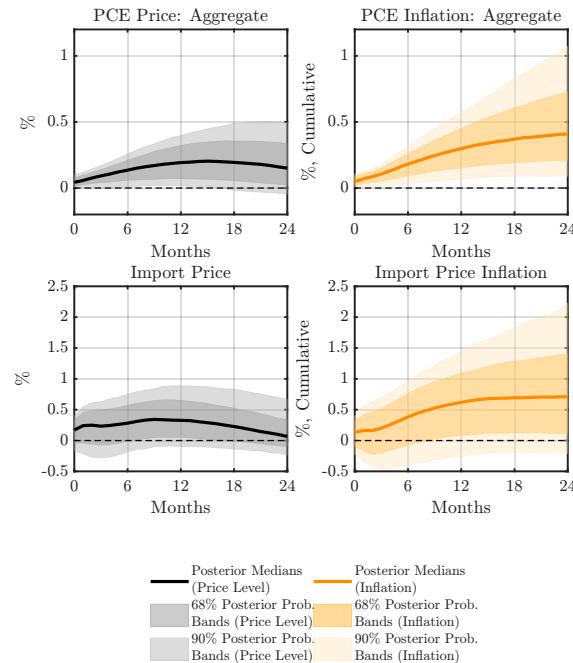


Figure G.6: Price-Level vs. Cumulated-Inflation IRFs: Adverse Supply Chain Shock

Notes. Price-level (gray) and cumulated inflation (orange) IRFs for PCE and import prices after an adverse supply chain shock.

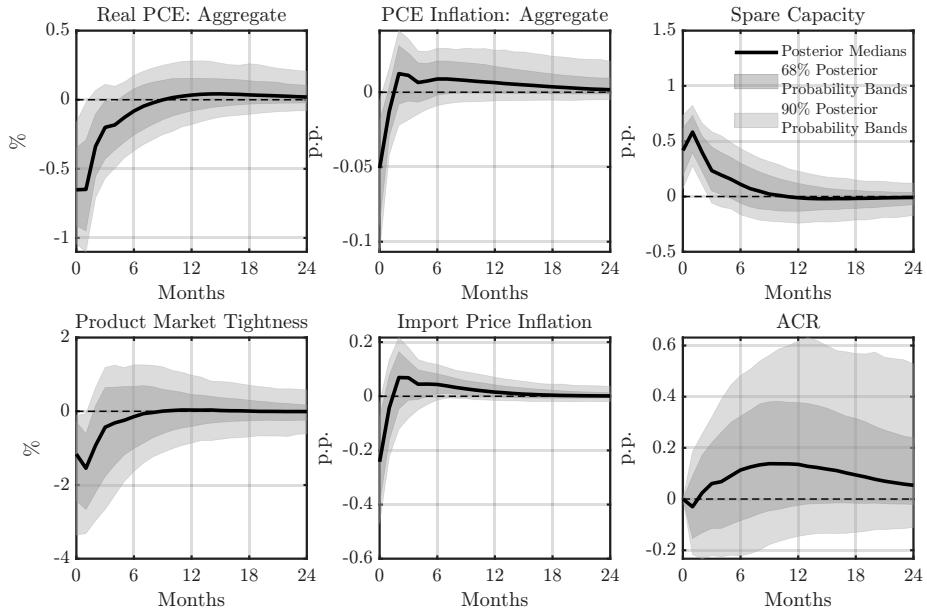


Figure G.7: IRFs to an Adverse Aggregate Demand Shock: PCE and Import Price Inflation

Notes. IRFs to a one-standard-deviation adverse aggregate demand shock using the baseline SVAR specification, except that the PCE chain-type price index and the import price index are replaced with their monthly inflation rates (log differences).

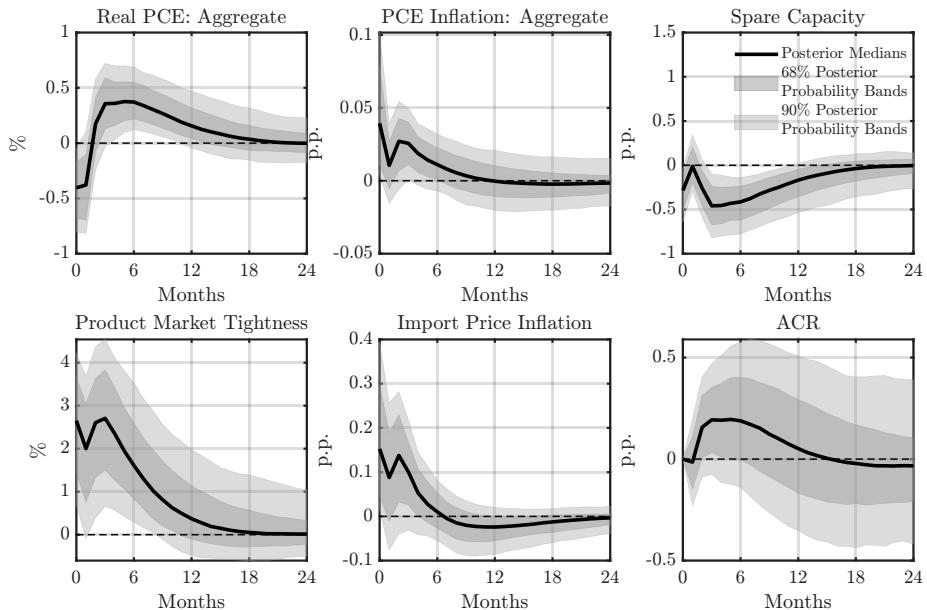


Figure G.8: IRFs to an Adverse Productive Capacity Shock: PCE and Import Price Inflation

Notes. IRFs to a one-standard-deviation adverse productive capacity shock using the baseline SVAR specification, except that the PCE chain-type price index and the import price index are replaced with their monthly inflation rates (log differences).

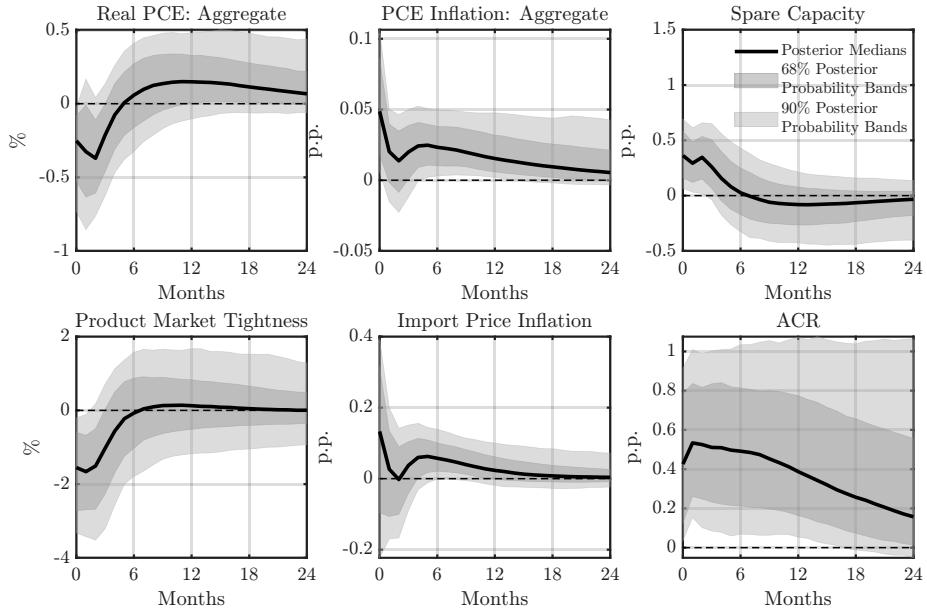


Figure G.9: IRFs to an Adverse Supply Chain Shock: PCE and Import Price Inflation

Notes. IRFs to a one-standard-deviation adverse supply chain shock using the baseline SVAR specification, except that the PCE chain-type price index and the import price index are replaced with their monthly inflation rates (log differences).

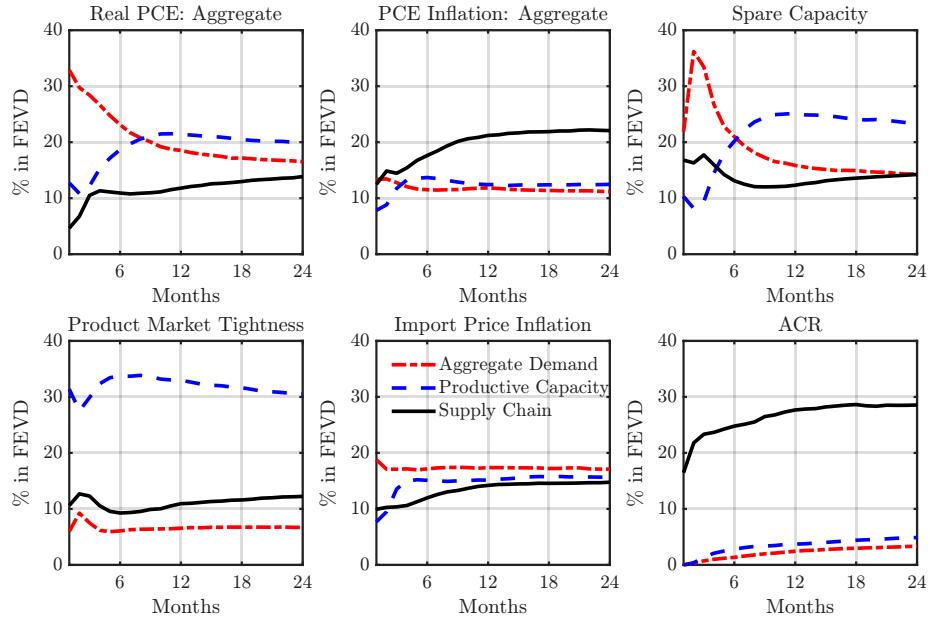


Figure G.10: FEVDs from the SVAR: PCE and Import Price Inflation

Notes. Each line reports the posterior-median share of the forecast error variance of an endogenous variable that is attributed to each of the three identified structural shocks across horizons. FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the PCE chain-type price index and the import price index are replaced with their monthly inflation rates (log differences).

G.3. Invertibility Under an Expanded Information Set

This appendix assesses whether the baseline six-variable SVAR omits forecast-relevant information that could undermine invertibility.¹⁴ In the SVAR context, non-invertibility arises when the model’s information set is too limited, so structural shocks cannot be expressed as linear combinations of past reduced-form innovations, and the identified shocks may conflate omitted disturbances (Fernández-Villaverde et al., 2007). To address this concern, we enlarge the information set to include additional macroeconomic and cost-push indicators and re-estimate the model, while maintaining the baseline identification on the original block.

The augmented specification adds five endogenous variables —real PCE of services, the PCE services price index, the West Texas Intermediate (WTI) spot price, the effective federal funds rate, and average hourly earnings— yielding an eleven-variable SVAR.¹⁵ Real services PCE and the PCE services price index capture potential reallocation between goods and services that may confound the aggregate demand block. The WTI spot price accounts for oil shocks that affect import costs and domestic prices. The federal funds rate controls for endogenous monetary policy responses. Average hourly earnings capture wage-driven cost-push pressures not fully reflected in price indices. Together, these additions create a richer information set that addresses the most likely sources of omitted forecast-relevant information.

We keep the sample, lag length, deterministic terms, and priors identical to the baseline and re-identify the same three shocks —aggregate demand, productive capacity, and supply chain— using the original sign and zero restrictions on the baseline block at horizon $k = 1$.¹⁶

Figures G.11–G.13 present the IRFs from the baseline and augmented SVARs. Three results stand out. First, for demand and capacity shocks, the IRFs remain largely invariant across specifications in both shape and magnitude, showing that the baseline six-variable system already spans the forecast-relevant information for these two shock blocks. Second, for the supply chain disturbance, the responses of the PCE and import price indices become smaller and less precisely estimated once services activity and prices, oil prices, the policy rate, and wages are added. Third,

¹⁴We thank an anonymous referee for suggesting this check.

¹⁵All additional series used in this appendix are retrieved from FRED with the following mnemonics: DGDSRX1 (real PCE of goods), PCESC96 (real PCE of services), DGDSRG3M086SBEA (PCE goods price index), DSERRG3M086SBEA (PCE services price index), WTISPLC (WTI spot price), FEDFUNDS (federal funds rate), and CES0500000003 (average hourly earnings). All series are seasonally adjusted except for the federal funds rate.

¹⁶All identification restrictions are imposed on the baseline variables; the added series are left unrestricted with respect to the three identified shocks.

the ACR behaves similarly across specifications for demand and capacity shocks, but its response to a supply chain shock is attenuated in the augmented system.

The attenuation of supply chain price responses is a mechanical result of expanding the information set: variation that the smaller system could allocate to the “supply chain” shock is reassigned to explicit cost-push and policy channels once additional variables enter the SVAR. Nonetheless, the augmented model supports the baseline six-variable SVAR results.

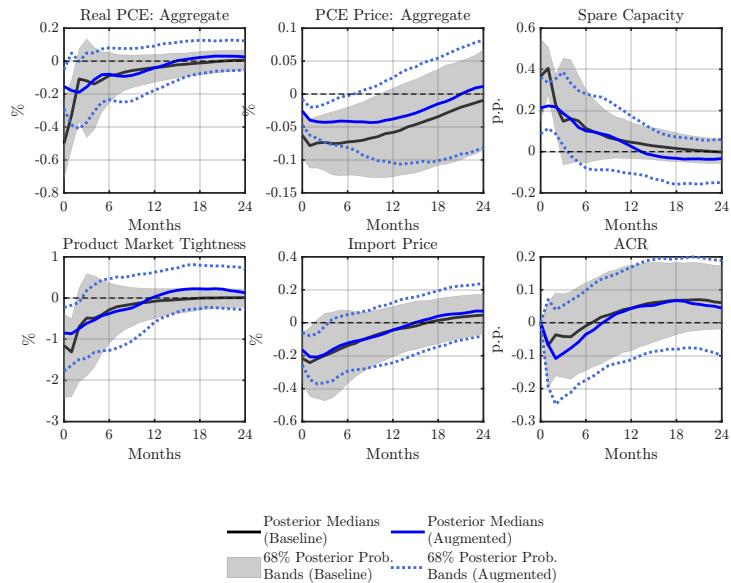


Figure G.11: IRFs to an Adverse Aggregate Demand Shock: Baseline vs. Augmented

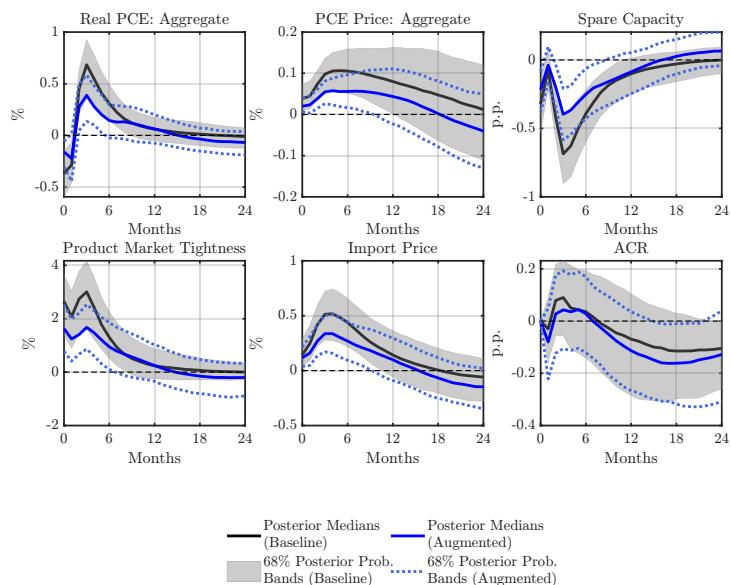


Figure G.12: IRFs to an Adverse Productive Capacity Shock: Baseline vs. Augmented

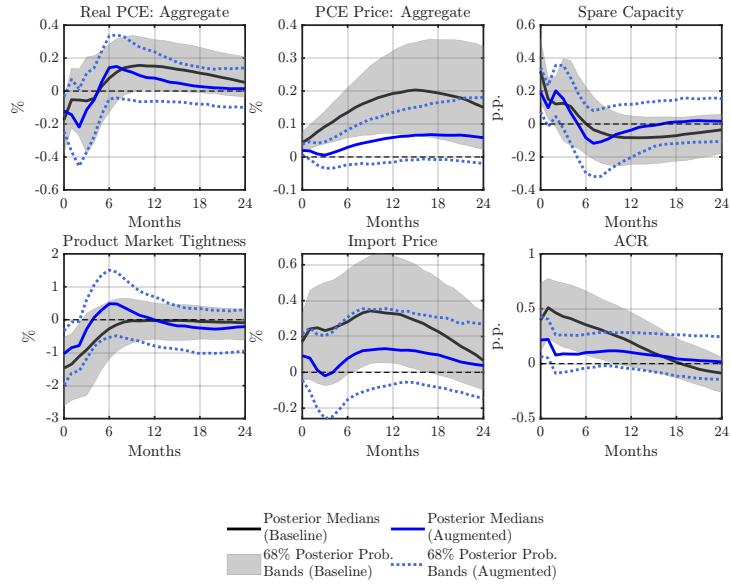


Figure G.13: IRFs to an Adverse Supply Chain Shock: Baseline vs. Augmented

Finally, as shown in Figures G.14–G.16, replacing real aggregate PCE and the PCE price index with their goods components (while holding the other variables fixed) and comparing this system to an augmented one that includes the services block, oil prices, the policy rate, and wages leads to the same broad conclusion. The augmented specification supports the conclusion that we have invertibility.

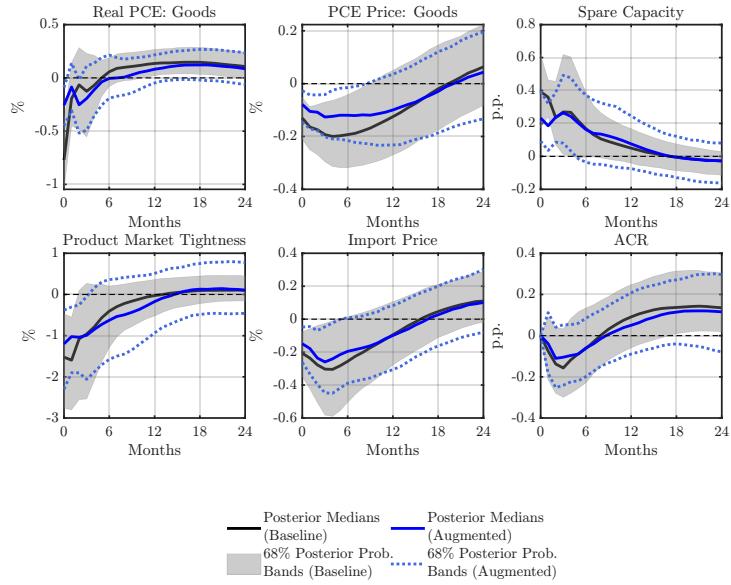


Figure G.14: IRFs to an Adverse Goods Demand Shock: Goods-Focused vs. Augmented

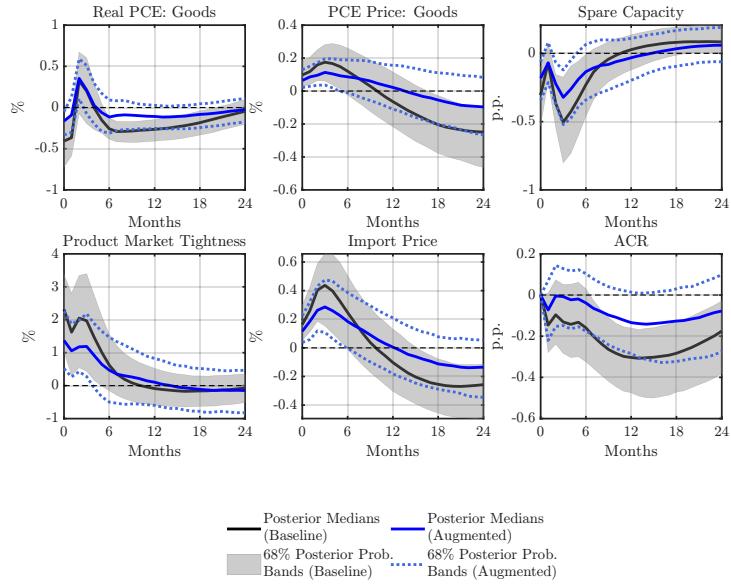


Figure G.15: IRFs to an Adverse Productive Capacity Shock: Goods-Focused vs. Augmented

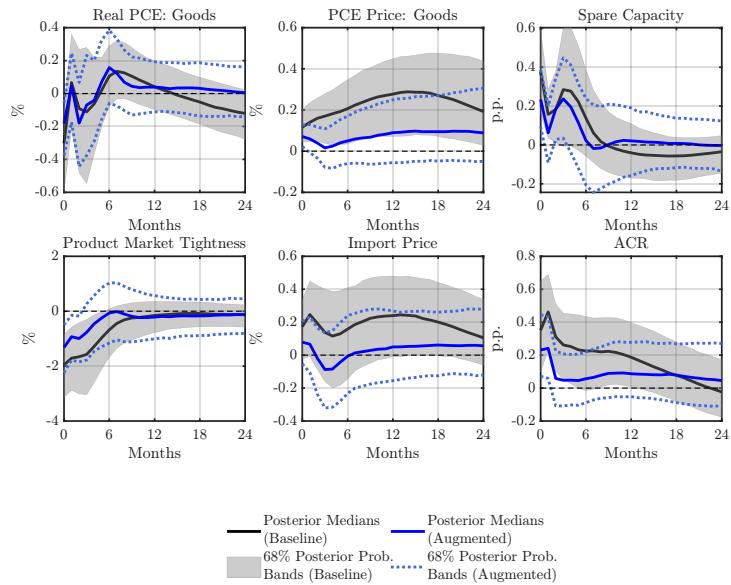


Figure G.16: IRFs to an Adverse Supply Chain Shock: Goods-Focused vs. Augmented

G.4. Nonlinear Transmission by Shock Size

Research on production networks (e.g., [Baaqee and Farhi, 2019, 2022](#)) shows that shock propagation can be nonlinear, particularly for large disturbances. To assess whether the transmission mechanisms of the three identified shocks—aggregate demand, productive capacity, and supply chain—depend on shock magnitude, we develop a nonlinear LP framework. This test is especially relevant here, given the intrinsic nonlinearities in aggregate demand and supply curves, as well

as in the match creation and separation conditions studied in Proposition 2 and Appendix E.3.

Let s_t denote the identified aggregate demand, productive capacity, or supply chain shock at time t . For each posterior draw $i = 1, \dots, N$ from the baseline Bayesian SVAR, we compute a draw-specific threshold $c^{(i)} = \text{median}(|s_t^{(i)}|)$ and define indicators:

$$\text{Small}_t^{(i)} = \mathbf{1}\left(|s_t^{(i)}| \leq c^{(i)}\right), \quad \text{Large}_t^{(i)} = 1 - \text{Small}_t^{(i)}.$$

For each horizon $k = 0, 1, \dots, K$, we estimate:

$$y_{t+k} = \alpha_k + \lambda_k t + \beta_k^{(i)} \left(s_t^{(i)} \text{Small}_t^{(i)} \right) + \gamma_k^{(i)} \left(s_t^{(i)} \text{Large}_t^{(i)} \right) + u_{k,t+k}^{(i)}, \quad (\text{G.1})$$

where y_{t+k} is the k -period-ahead value of an endogenous variable in the Baseline SVAR, α_k is a horizon-specific constant, t is a deterministic time trend, and $u_{k,t+k}^{(i)}$ is the k -step-ahead forecast error for draw i . To reduce noise at longer horizons, we estimate Equation (G.1) using the SLP method (Barnichon and Brownlees, 2019), applying Newey–West standard errors (Newey and West, 1987) to obtain inference for $\beta_k^{(i)}$ and $\gamma_k^{(i)}$.¹⁷

To aid interpretation, we scale coefficients by the within-draw conditional standard deviations:

$$\sigma_{\text{small}}^{(i)} = \text{sd}\left(s_t^{(i)} \mid \text{Small}_t^{(i)} = 1\right), \quad \sigma_{\text{large}}^{(i)} = \text{sd}\left(s_t^{(i)} \mid \text{Large}_t^{(i)} = 1\right).$$

The shock-size-specific IRFs are:

$$\text{IRF}_{\text{small}}^{(i)}(k) = \beta_k^{(i)} \sigma_{\text{small}}^{(i)}, \quad \text{IRF}_{\text{large}}^{(i)}(k) = \gamma_k^{(i)} \sigma_{\text{large}}^{(i)}.$$

We report pointwise posterior medians with 68% credible intervals across i for each k .

This procedure tests whether large shocks propagate differently from small ones, as expected under nonlinear network amplification, while incorporating uncertainty in shock identification. Because the sample includes the COVID-19 pandemic, which plausibly represents the largest shock in the period, this analysis ensures that the documented dynamics are not driven by a single exceptional episode.

Figures G.17–G.19 plot the IRFs to adverse aggregate demand, productive capacity, and supply chain shocks, respectively, with each shock split into large and small categories. The median responses to large shocks differ in both magnitude and profile from those to small shocks, but the posterior probability bands overlap across most horizons. We read this as evidence that

¹⁷See Appendix B for details on SLP.

nonlinearities may matter, especially for large shocks, while our main conclusions remain intact: the transmission mechanism is qualitatively stable across shock sizes, and the macroeconomic effects of the identified shocks are not driven by moderate nonlinear amplification.

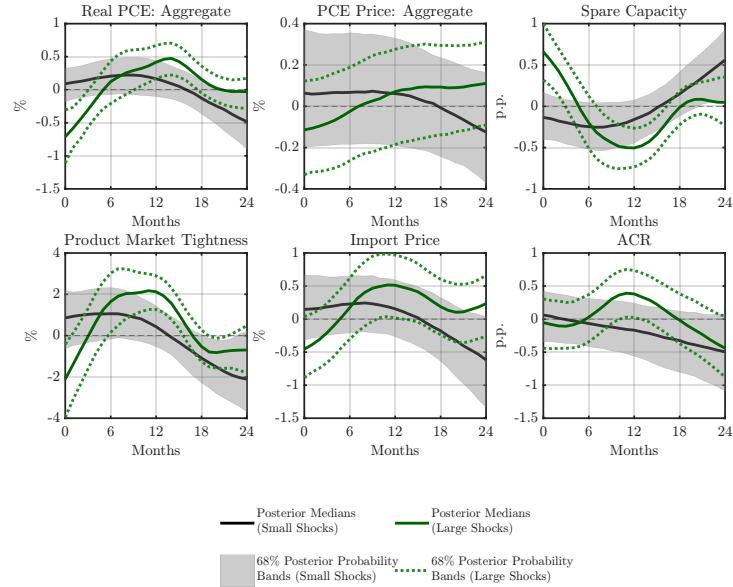


Figure G.17: Nonlinear Transmission of Aggregate Demand Shocks by Size

Notes. IRFs to small and large adverse aggregate demand shocks, with shock size classified by the median magnitude of the identified shock series from the baseline SVAR.

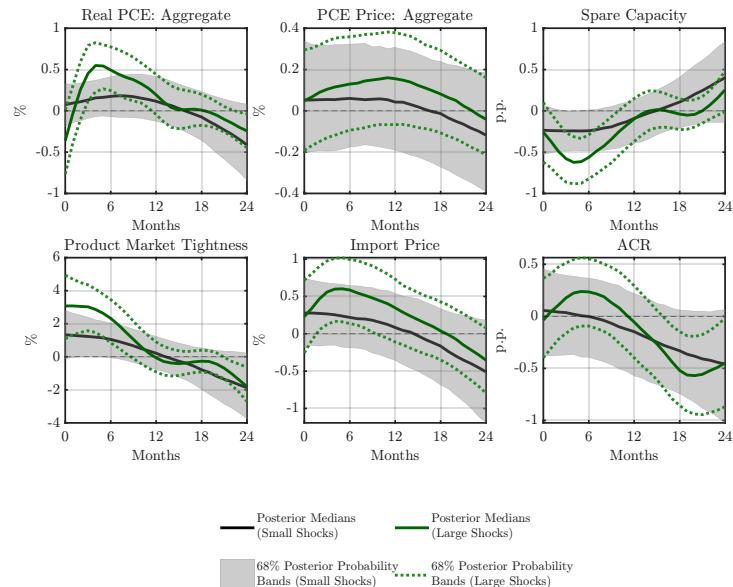


Figure G.18: Nonlinear Transmission of Productive Capacity Shocks by Size

Notes. IRFs to small and large adverse productive capacity shocks, with shock size classified by the median magnitude of the identified shocks from the baseline SVAR.

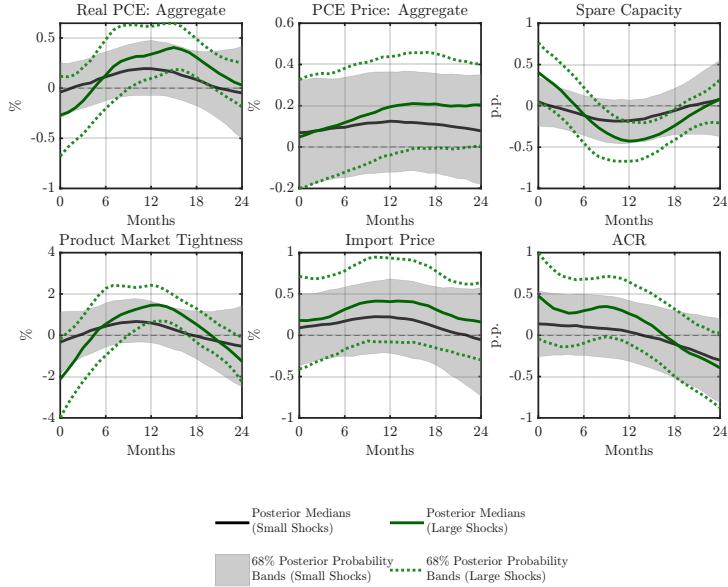


Figure G.19: Nonlinear Transmission of Supply Chain Shocks by Size

Notes. IRFs to small and large adverse supply chain shocks, with shock size classified by the median magnitude of the identified shocks from the baseline SVAR.

G.5. Different Lag Structures

To examine whether our results depend on the number of lags, we re-estimate the baseline SVAR with four and six lags, holding the sample, priors, identification, and other settings fixed.

Three findings stand out. First, median IRFs are highly similar across lag specifications: as shown in Figures G.20–G.25, there are no sign reversals, shifts in peak timing, or changes in overall dynamics. Second, uncertainty bands widen only modestly with more lags, mainly in the first 6–8 months; at medium horizons they overlap substantially, and long-horizon differences remain minor. Third, the key narratives discussed in Section 4 hold unchanged: an adverse supply chain disturbance still dampens real activity and product market tightness, raises spare capacity, and pushes prices upward.

We retain two lags in the baseline specification for parsimony and robustness. With monthly data and a relatively small SVAR, additional lags quickly enlarge the parameter space and reduce estimation efficiency without adding empirical content. Since IRFs and qualitative inference are insensitive to lag length, $L = 2$ captures the essential dynamics while avoiding overfitting.

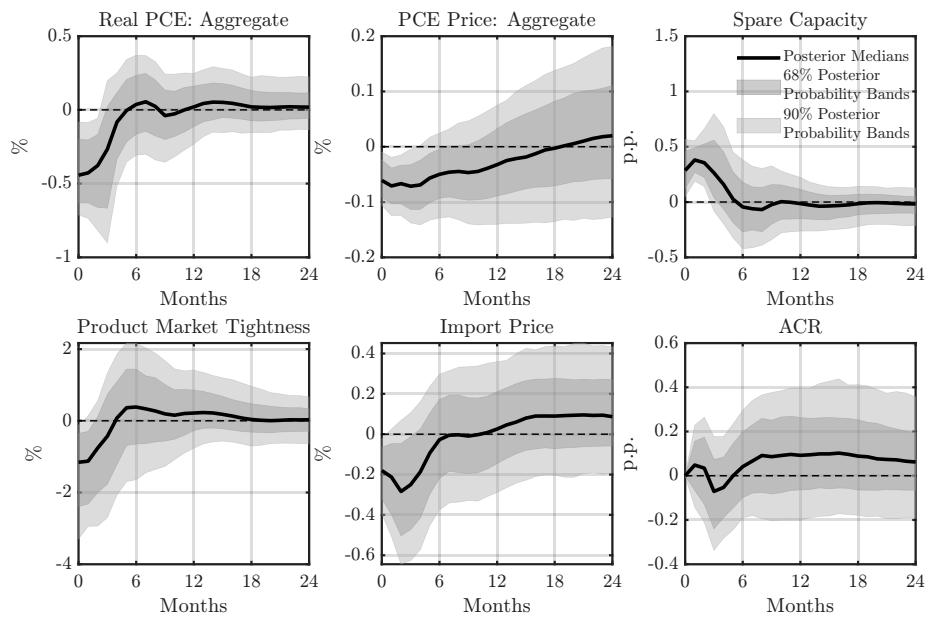


Figure G.20: IRFs to an Adverse Aggregate Demand Shock: Four Lags

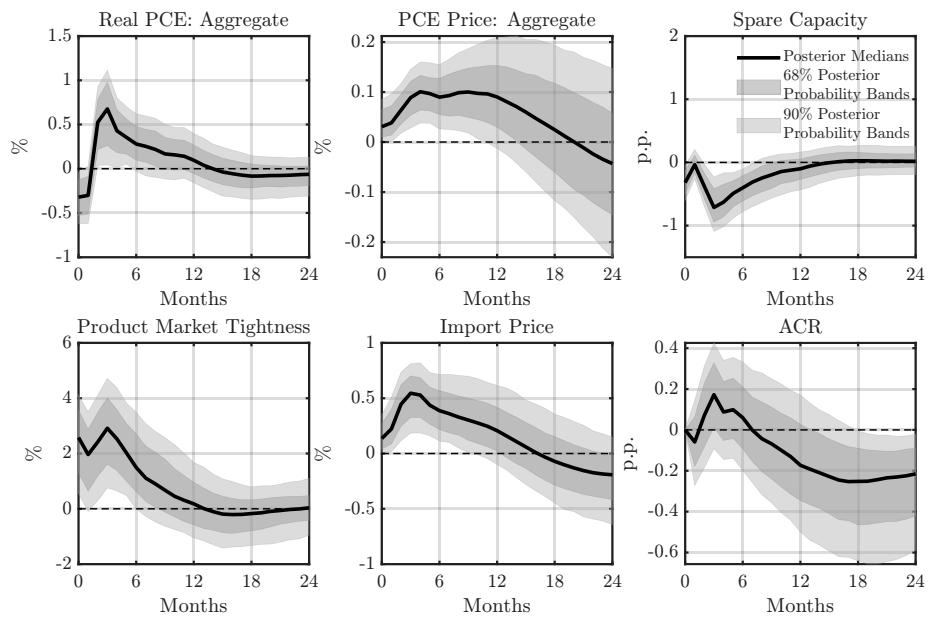


Figure G.21: IRFs to an Adverse Productive Capacity Shock: Four Lags

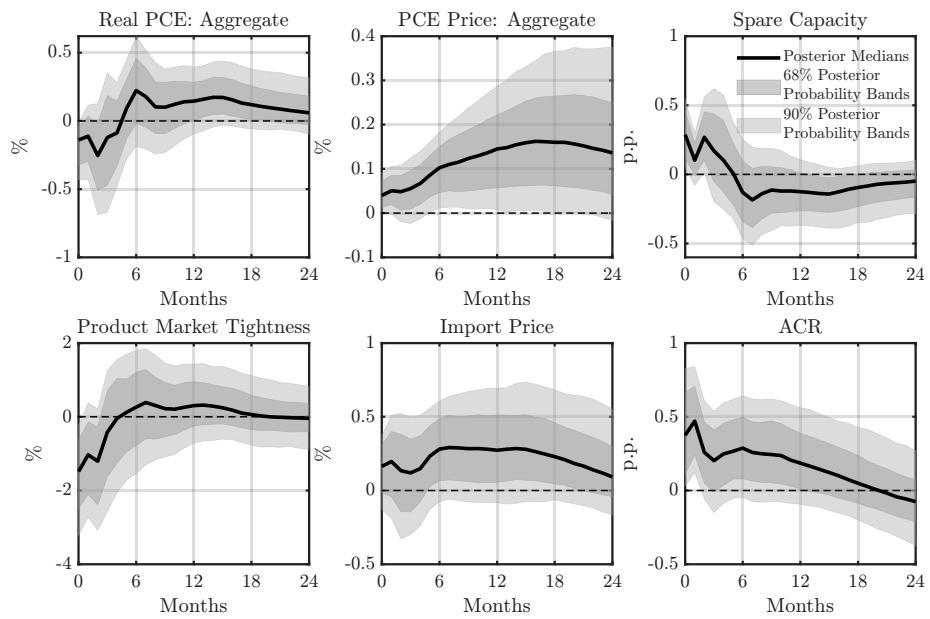


Figure G.22: IRFs to an Adverse Supply Chain Shock: Four Lags

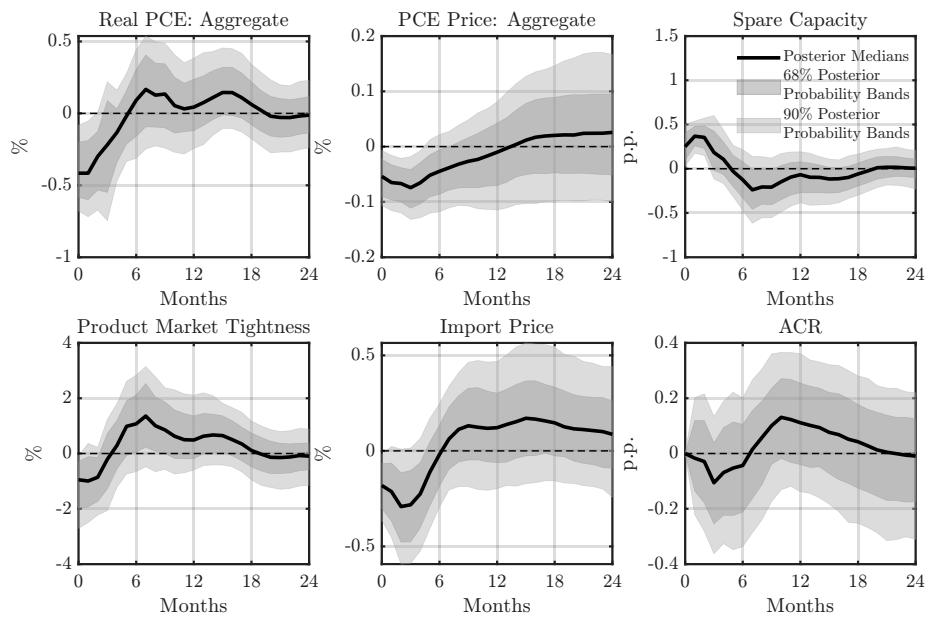


Figure G.23: IRFs to an Adverse Aggregate Demand Shock: Six Lags

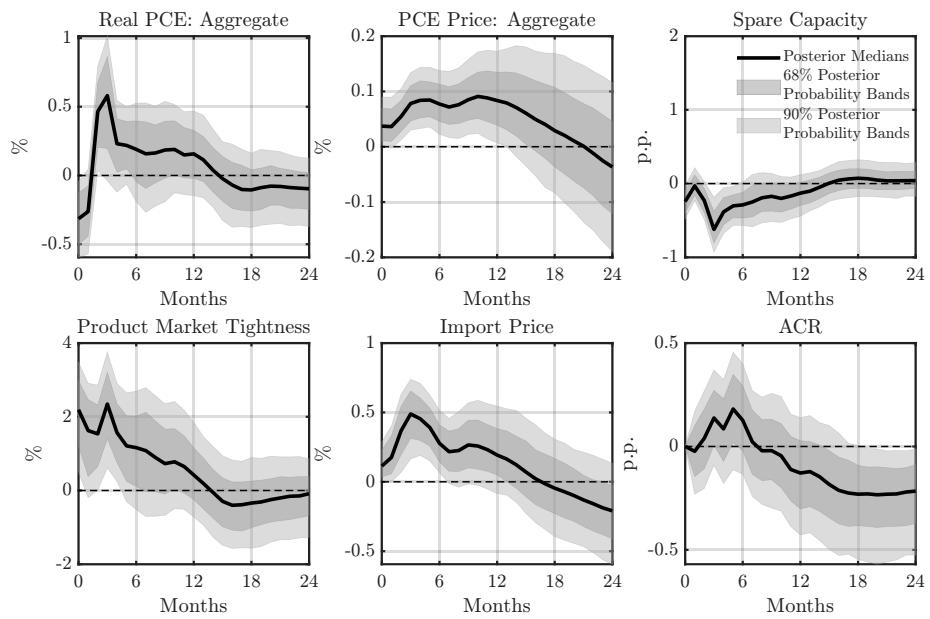


Figure G.24: IRFs to an Adverse Productive Capacity Shock: Six Lags

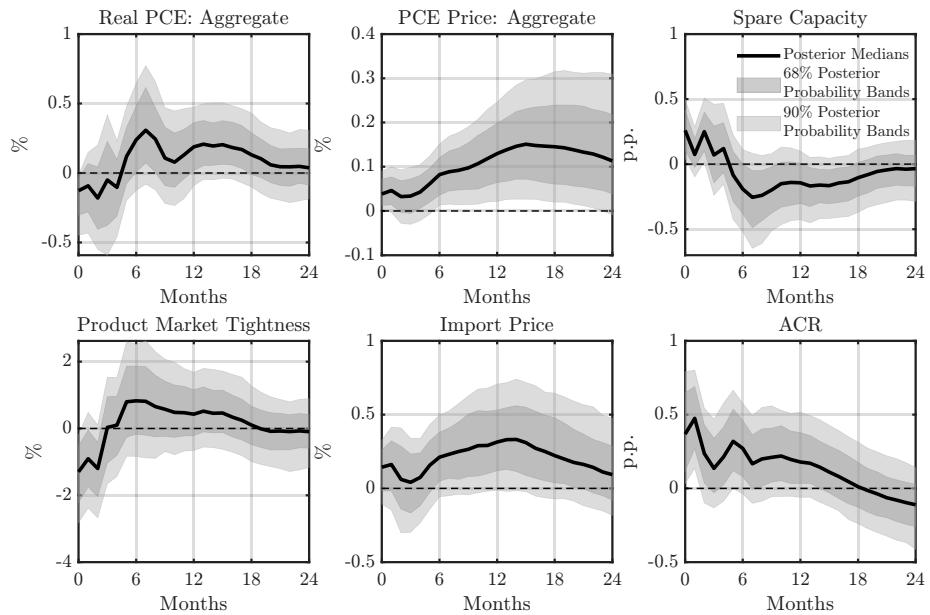


Figure G.25: IRFs to an Adverse Supply Chain Shock: Six Lags

G.6. Dropping the Linear Trend

We re-estimate the baseline SVAR with an intercept only, holding the sample, lag length, priors, shock normalization, and identification (sign/zero restrictions) fixed. This checks whether allowing for a linear drift in level variables affects impulse responses or variance shares.

Figures G.26–G.28 show the IRFs under the constant-only specification. For the demand shock, the results closely track the baseline: real PCE falls on impact and mean reverts; spare capacity rises and then fades; the price level dips slightly before returning toward zero; and the responses of product market tightness, import prices, and ACR remain small and short lived. Posterior probability bands coincide across the two specifications, and any level shifts are minor.

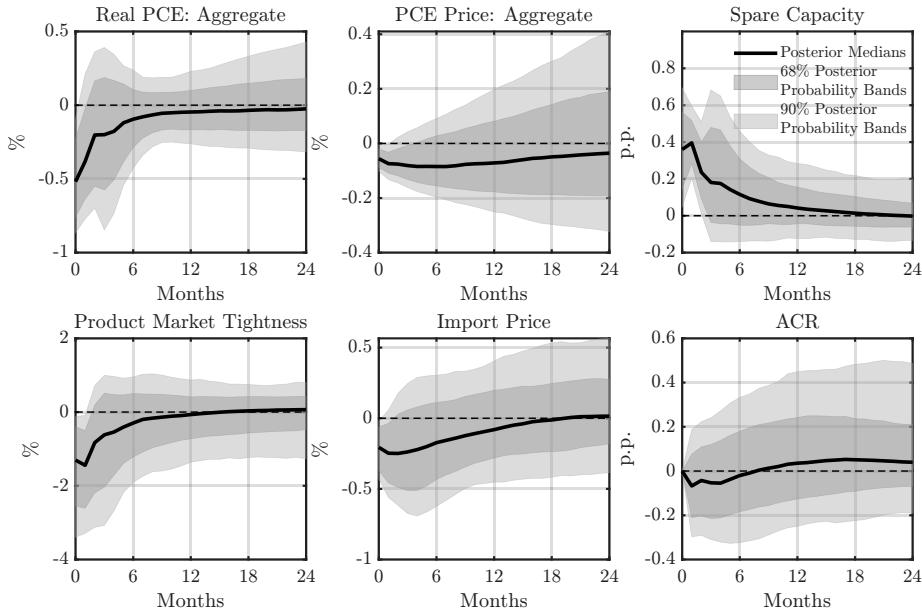


Figure G.26: IRFs to an Adverse Aggregate Demand Shock: Constant Only

For the capacity shock, activity responses retain their familiar hump-shaped dynamics with gradual reversion. The price-level response is nearly identical at short and medium horizons; if anything, removing the trend induces a slightly more negative drift at long horizons, but the change remains within posterior uncertainty.

For the supply chain shock, the qualitative pattern is unchanged across specifications: the price level rises on impact and stays elevated at medium horizons; real activity eases, and spare capacity increases; import prices and ACR jump contemporaneously and then recede. Relative to the baseline, the constant-only specification delivers very similar peak timing and magnitudes for real and nominal variables, with broadly overlapping uncertainty bands. If there is any detectable

change, it is a mild increase in medium-horizon persistence of the price-level response, but its shape and economic interpretation stay the same.

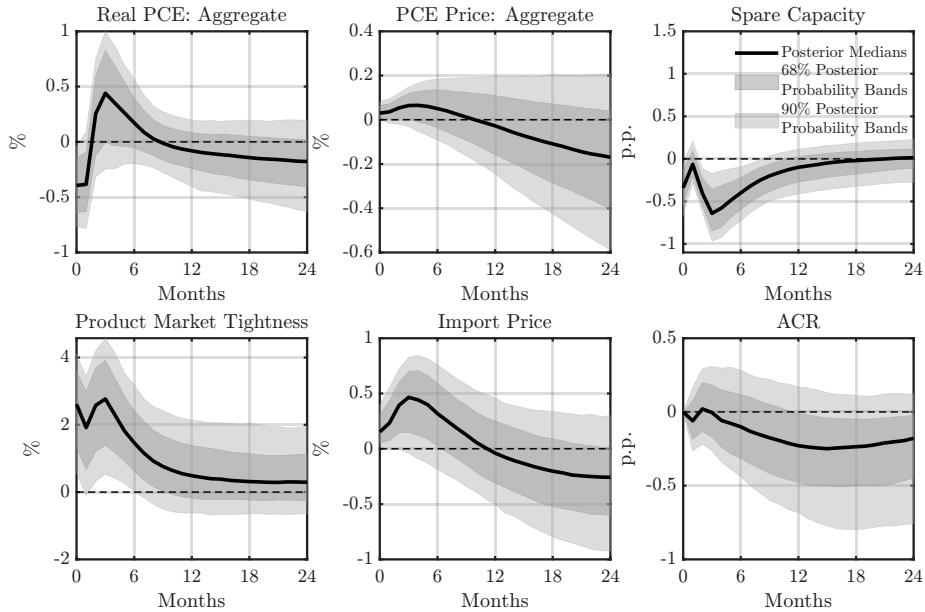


Figure G.27: IRFs to an Adverse Productive Capacity Shock: Constant Only

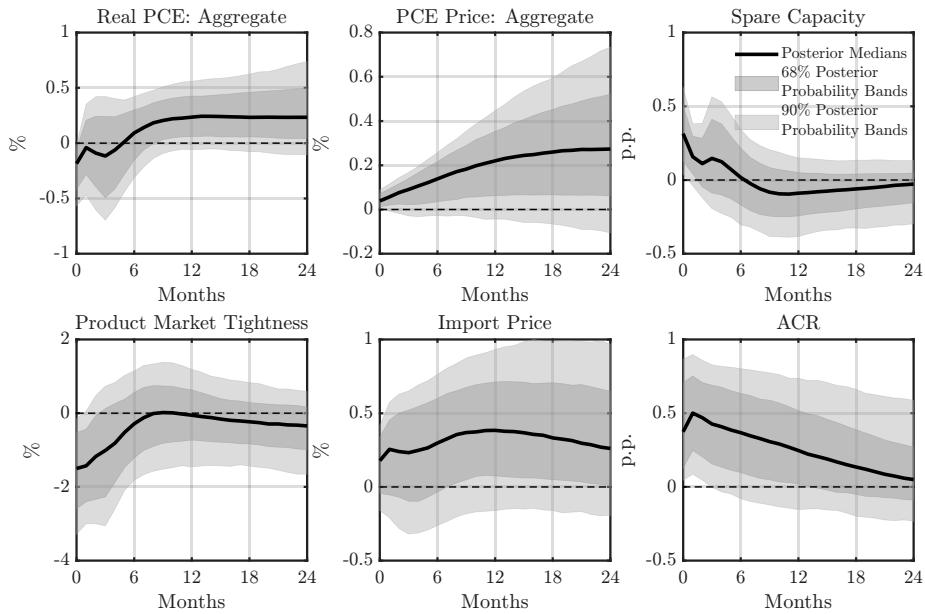


Figure G.28: IRFs to an Adverse Supply Chain Shock: Constant Only

As shown in Figure G.29, the FEVDs remain stable across specifications. Supply chain shocks continue to be the dominant source of PCE price variation at medium horizons, while capacity shocks still account for most of the variance in spare capacity and product market tightness.

Dropping the trend shifts a modest share of variance toward supply chain shocks for real PCE and import prices at longer horizons, but the ranking of contributions is unchanged.

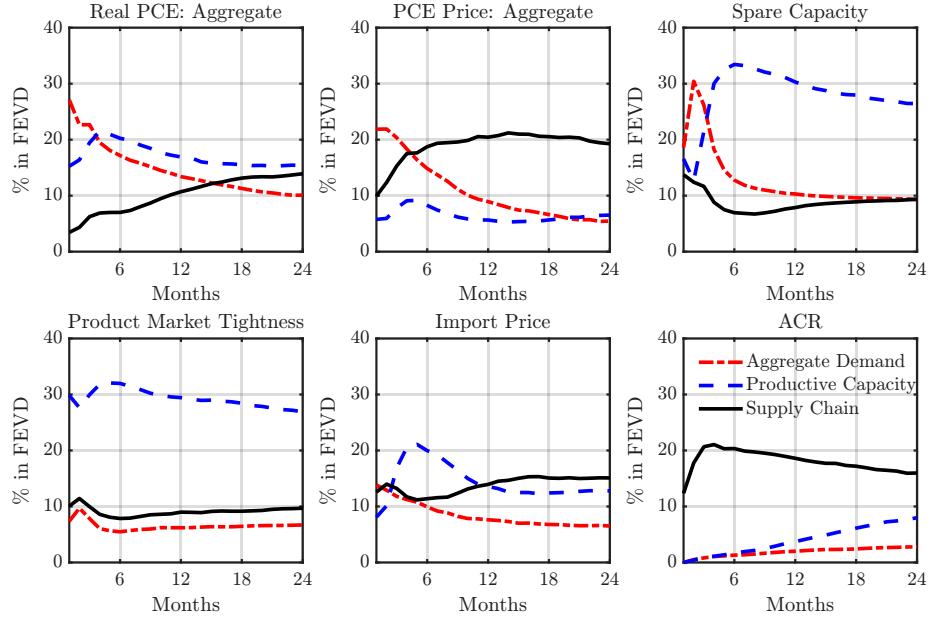


Figure G.29: FEVDs from the SVAR: Constant Only

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the linear time trend is replaced with an intercept only.

Taken together, allowing for a linear deterministic trend has limited influence on our conclusions. Because prices are modeled in levels, we retain the baseline with a linear trend to absorb secular drift and treat the constant-only specification as a robustness check that confirms the main narratives (demand and capacity dynamics and the prominent role of supply chain disturbances for inflation) are effectively unchanged.

G.7. Alternative Proxies for Activity and Prices

This appendix evaluates robustness to alternative measures of activity and prices. We consider three alternatives: (i) real PCE of goods and the PCE goods price index (FRED mnemonics `DGDSRX1`, `DGDSRG3M086SBEA`); (ii) IP and the PPI for final-demand finished goods (`INDPRO`, `WPSFD49207`); and (iii) real GDP and the GDP price deflator (`GDPC1`, `GDPDEF`). For (iii), the monthly real GDP series is constructed via Chow–Lin interpolation of the quarterly series using IP as the monthly indicator, while the monthly GDP deflator is obtained via Chow–Lin interpolation of the quarterly series using both CPI (`CPIAUCSL`) and PPI (`WPSFD49207`) as indicators,

following [Chow and Lin \(1971\)](#). All series are seasonally adjusted.

Goods PCE and prices. Replacing aggregate PCE and its chain-type price index with goods-only measures leaves the dynamics essentially unchanged. Goods demand shocks resemble the baseline: activity falls on impact and prices dip slightly before returning toward zero (Figure [G.30](#)). Production capacity and supply chain shocks are very similar to the baseline case (Figures [G.31-G.32](#)).

The FEVDs likewise preserve the baseline ranking: supply chain shocks account for a material share of price variance at medium horizons, while capacity shocks dominate activity variance after the first few months (Figure [G.33](#)). The HD in Figure [G.34](#) also highlights the drivers of PCE goods inflation discussed in Section 4, with the exception that productive capacity shocks played only a minimal role in pushing up inflation from late 2020 to mid 2021.

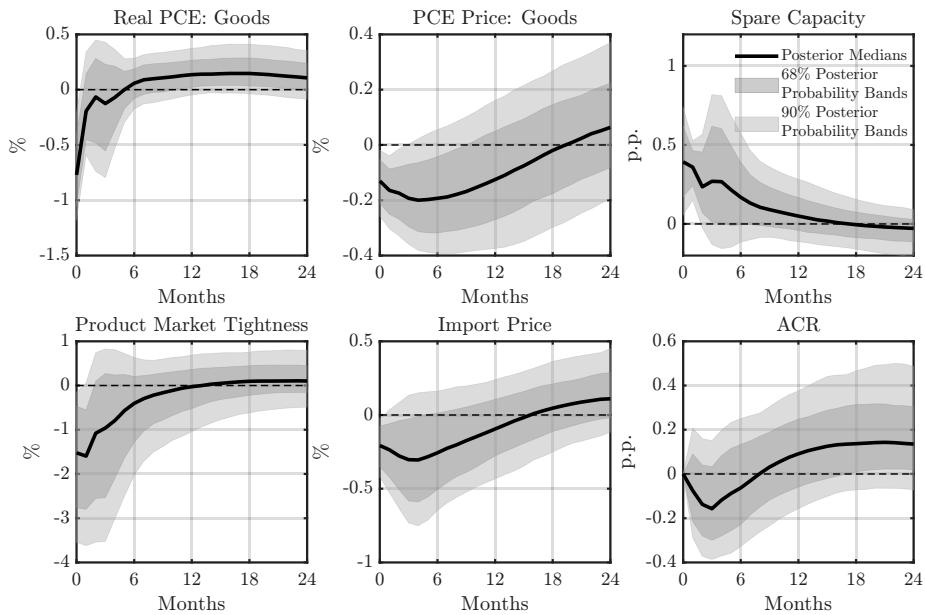


Figure G.30: IRFs to an Adverse Goods Demand Shock: Goods PCE and Prices

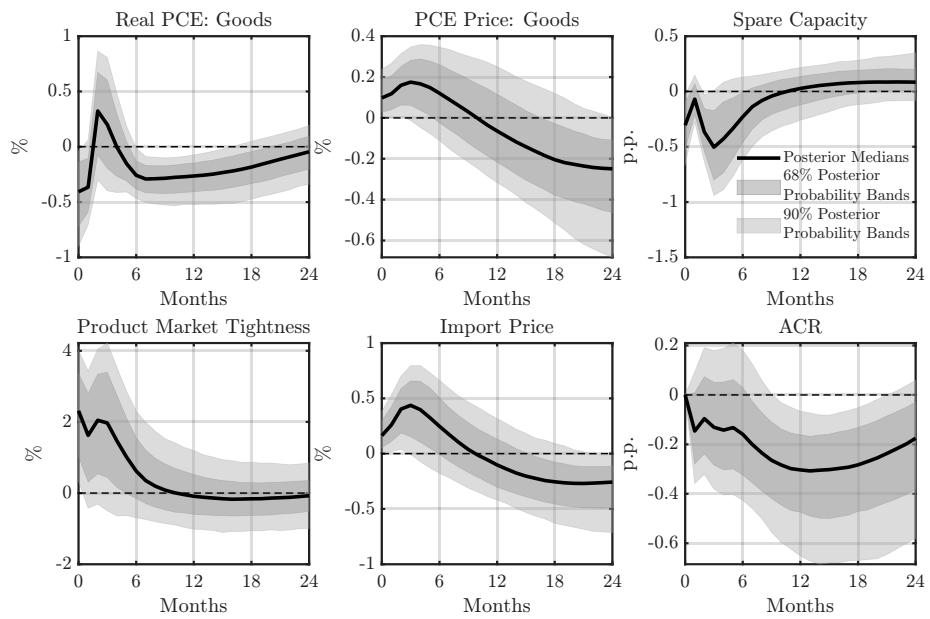


Figure G.31: IRFs to an Adverse Productive Capacity Shock: Goods PCE and Prices

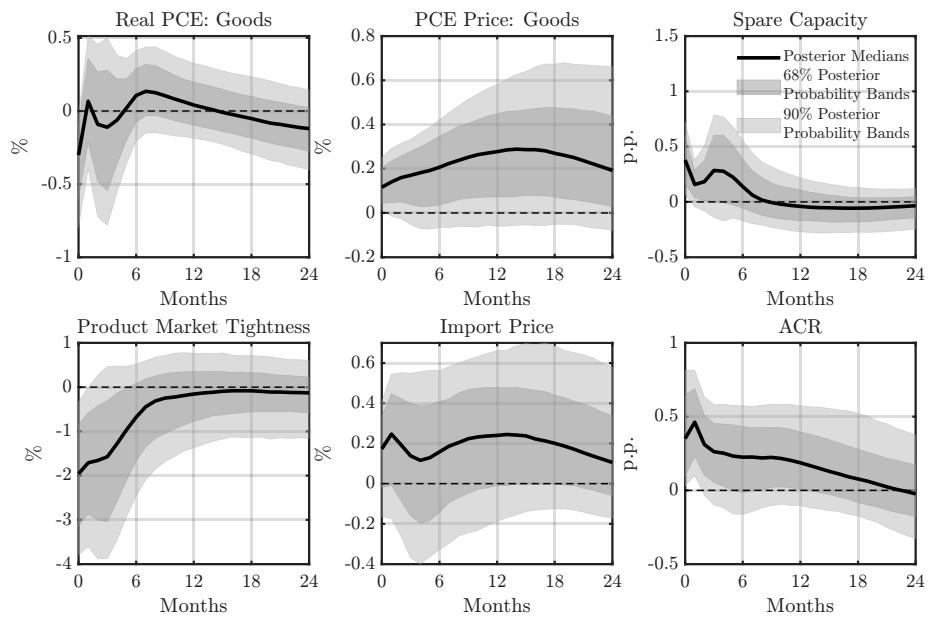


Figure G.32: IRFs to an Adverse Supply Chain Shock: Goods PCE and Prices

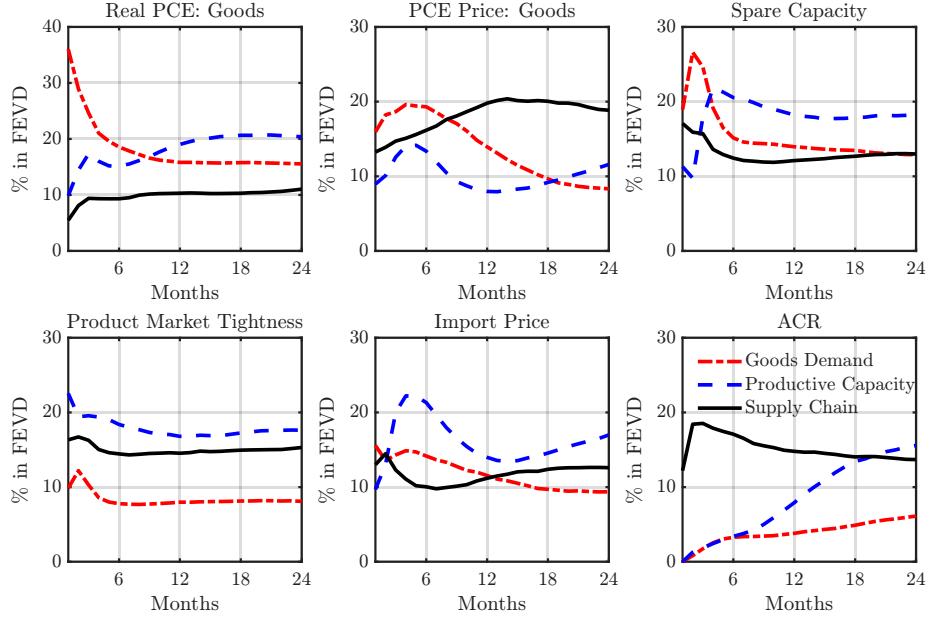


Figure G.33: FEVDs from the SVAR: Goods PCE and Prices

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that real PCE and the PCE price index are replaced by goods PCE and the PCE goods price index.

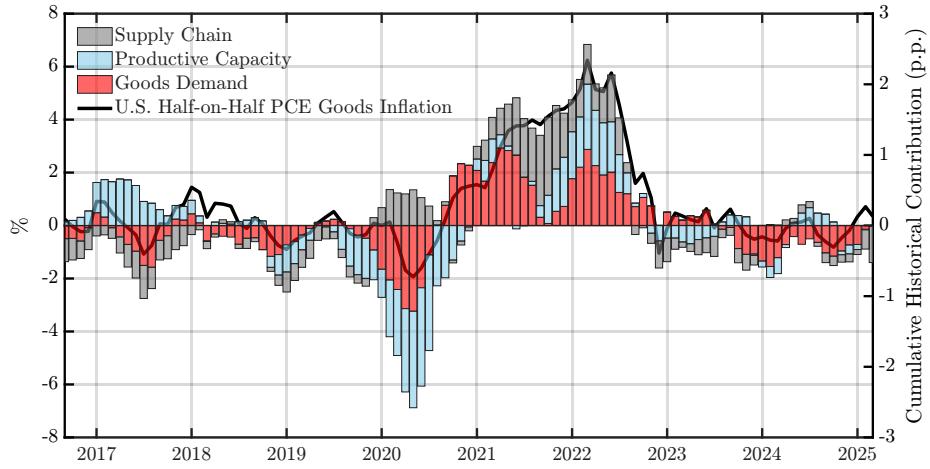


Figure G.34: HD of U.S. HoH PCE Goods Inflation

Notes. The solid line plots U.S. PCE goods inflation, measured as the HoH growth rate of the PCE goods price index, against the left axis (%). The shaded bars report the cumulative historical contributions of goods demand, productive capacity, and supply chain shocks, plotted against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the baseline specification, except that real PCE and the PCE price index are replaced with goods PCE and the PCE goods price index.

IP and PPI. Using IP and the finished-goods PPI slightly amplifies the price pass-through of a supply chain disturbance relative to the baseline: the PPI rises more and for longer (Figure G.37). Responses to demand and capacity shocks otherwise align with the baseline. The demand shock reduces activity on impact and induces, at most, a mild and short-lived disinflation. The capacity shock lowers activity, reduces spare capacity, and triggers a price jump on impact; in subsequent months, activity rebounds, spare capacity briefly retraces before declining again, and inflationary pressures continue to build (Figures G.35 and G.36).

The FEVDs reveal a slightly different pattern: while supply chain shocks still explain a material share of PPI variation across horizons, capacity shocks account for most of the medium-horizon variation in both activity and prices (Figure G.38), with demand shocks remaining concentrated at short horizons. Consistently, Figure G.39 shows that capacity shocks are the predominant driver of producer price inflation in the post-pandemic period, while demand and supply chain shocks play only minor roles.

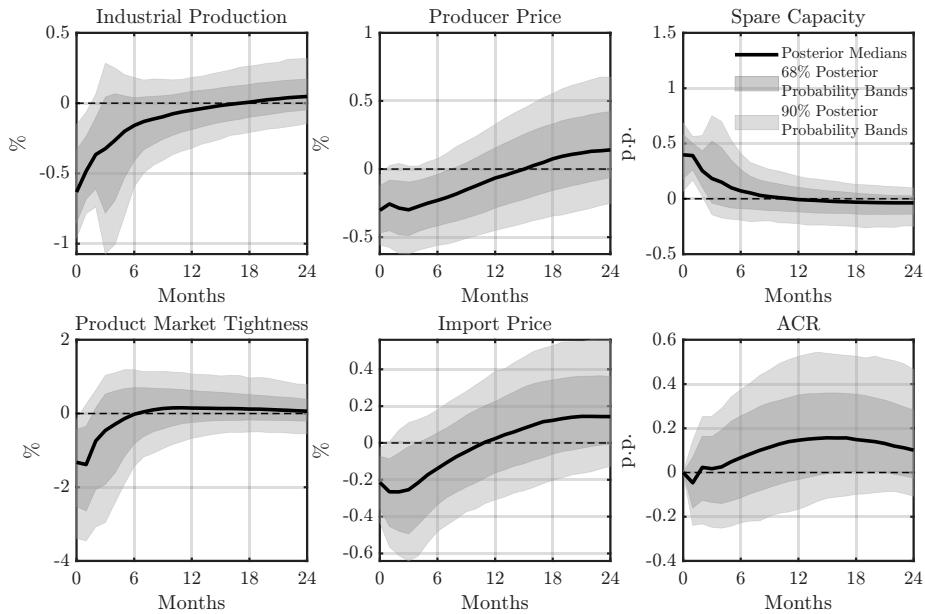


Figure G.35: IRFs to an Adverse Goods Demand Shock: IP and PPI

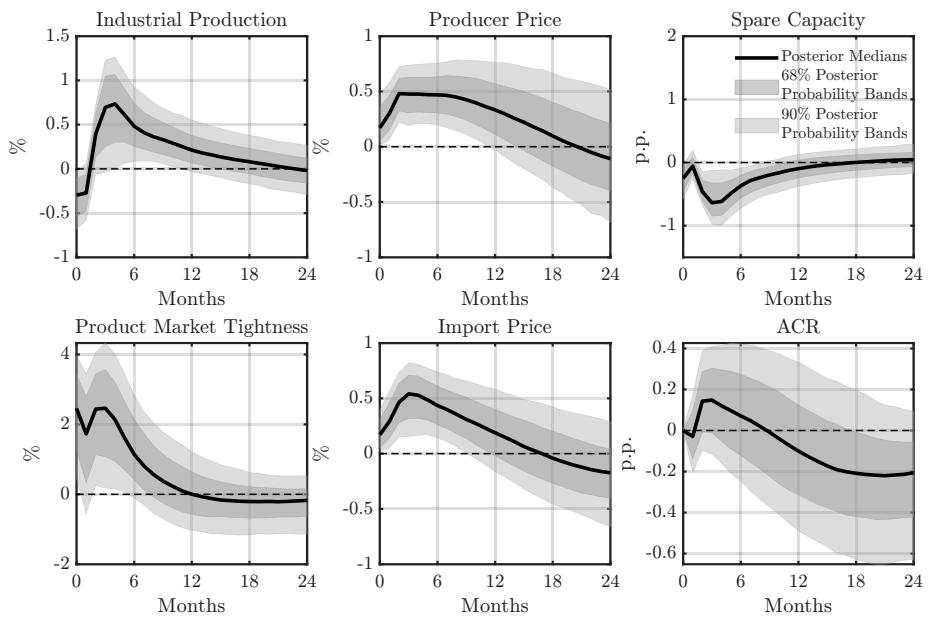


Figure G.36: IRFs to an Adverse Productive Capacity Shock: IP and PPI

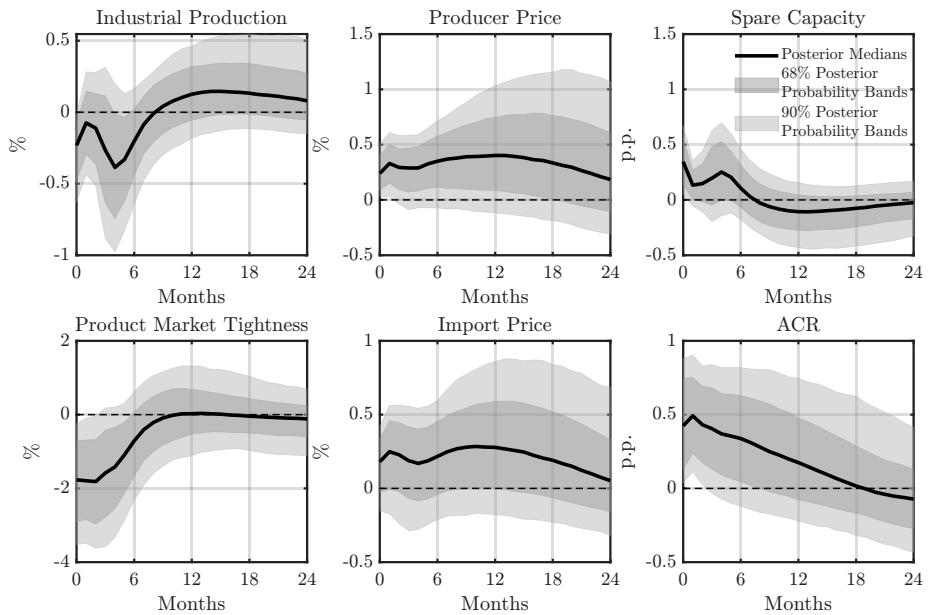


Figure G.37: IRFs to an Adverse Supply Chain Shock: IP and PPI

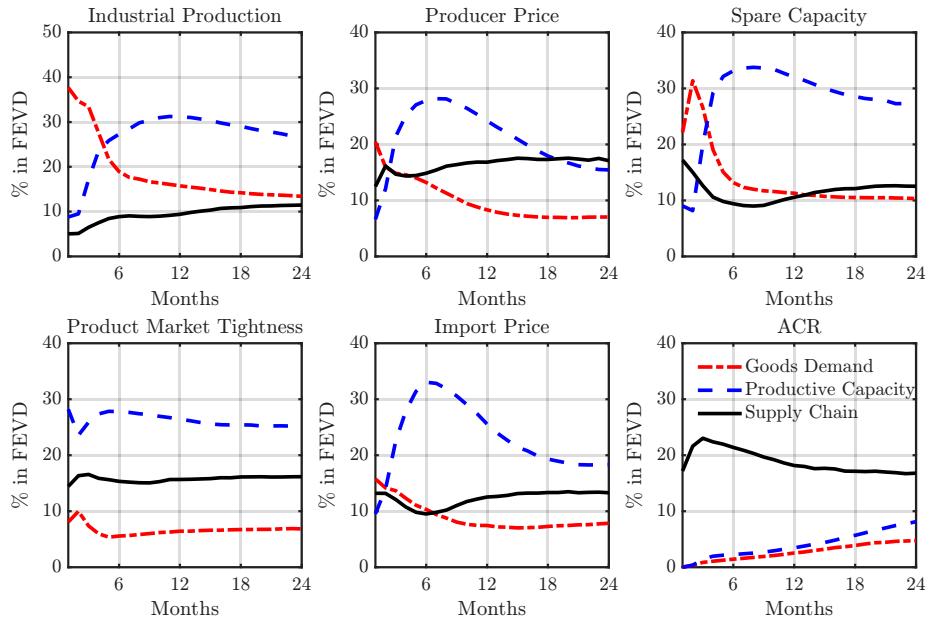


Figure G.38: FEVDs from the SVAR: IP and PPI

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that real PCE and the PCE price index are replaced by IP and the PPI.

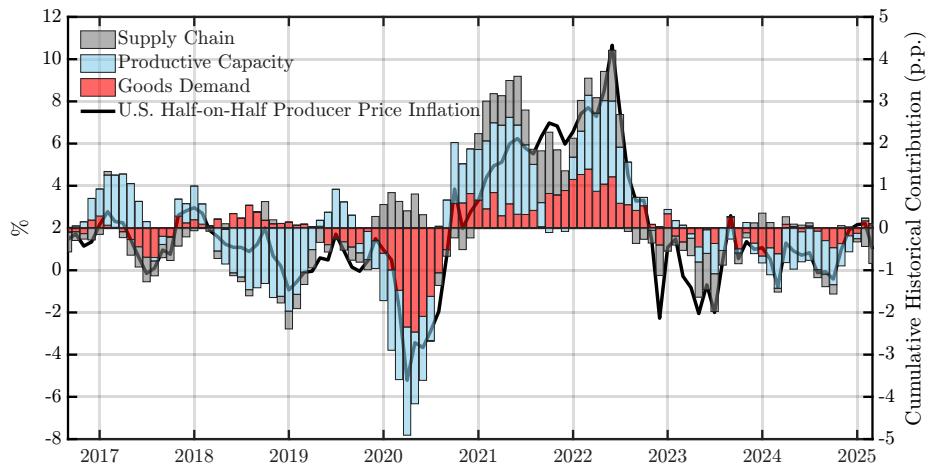


Figure G.39: HD of U.S. HoH Producer Price Inflation

Notes. The solid line plots U.S. producer price inflation, measured as the HoH growth rate of the PPI for final-demand finished goods, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of goods demand, productive capacity, and supply chain shocks, plotted against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that real PCE and the PCE price index are replaced by IP and the PPI.

Real GDP and GDP deflator. When we use the Chow–Lin–interpolated GDP and deflator, the median IRFs track the baseline closely but display wider posterior bands, reflecting the extra measurement uncertainty introduced by temporal disaggregation (Figures G.40–G.42). Demand shocks lower activity on impact, while the deflator response is attenuated and gradually mean reverts at longer horizons. Capacity shocks reproduce the familiar real-side dynamics and associated price responses, and supply chain shocks raise the deflator with hump-shaped pass through. The FEVDs preserve the baseline ranking: capacity shocks dominate real activity at medium horizons, supply chain shocks account for most price variance, and demand shocks remain concentrated at short horizons (Figure G.43). The HD in Figure G.44 likewise shows no major differences from the attribution in Figure 11.

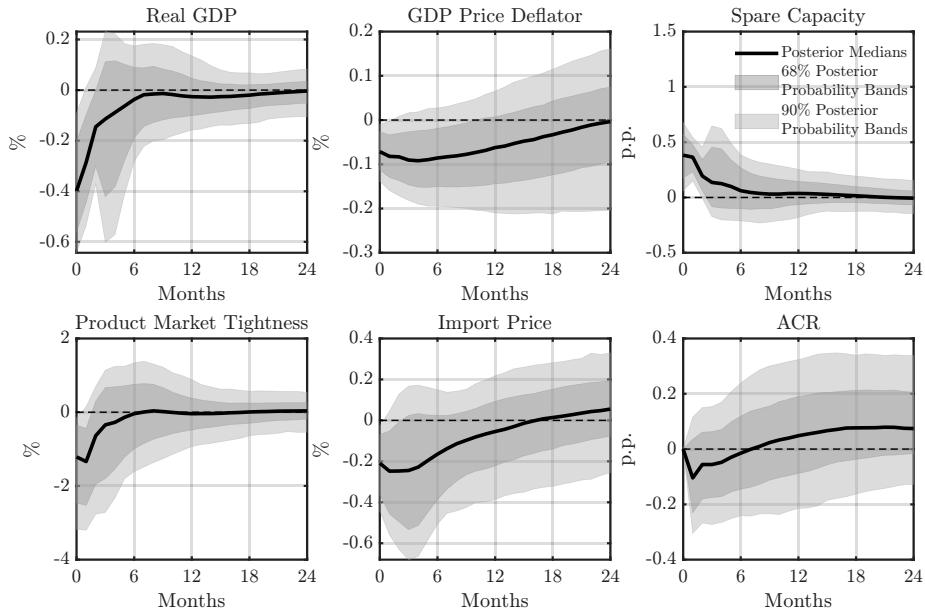


Figure G.40: IRFs to an Adverse Aggregate Demand Shock: Real GDP and GDP Deflator

Taking stock. In sum, four observations stand out. First, the macroeconomic impacts of demand and capacity shocks are robust across alternative measures of output and prices. Second, supply chain price effects are at least as strong, and often stronger, when producer- or GDP-based price indices are used, underscoring the role of logistics bottlenecks as an independent source of price variation. Third, the FEVDs consistently show that (i) capacity shocks dominate real activity at medium horizons, (ii) supply chain shocks account for a sizable share of price variance, and (iii) demand shocks remain concentrated at short horizons. This pattern holds in both the baseline FEVDs and all three alternatives. Finally, across different proxies for inflation, a large

share of the post-pandemic surge in U.S. inflation can be attributed to supply-side disruptions, with the relative importance of capacity versus supply chain shocks varying by measure.

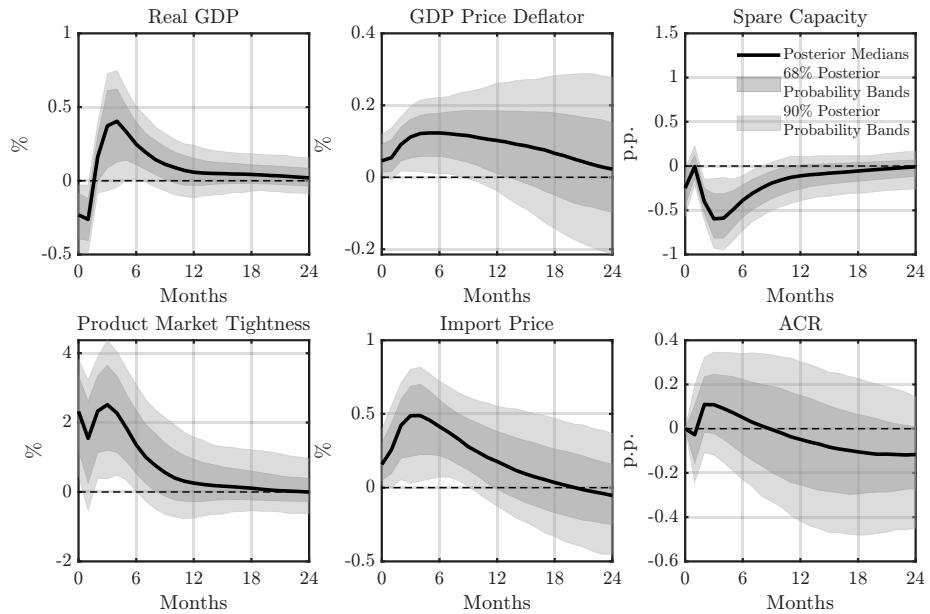


Figure G.41: IRFs to an Adverse Productive Capacity Shock: Real GDP and GDP Deflator

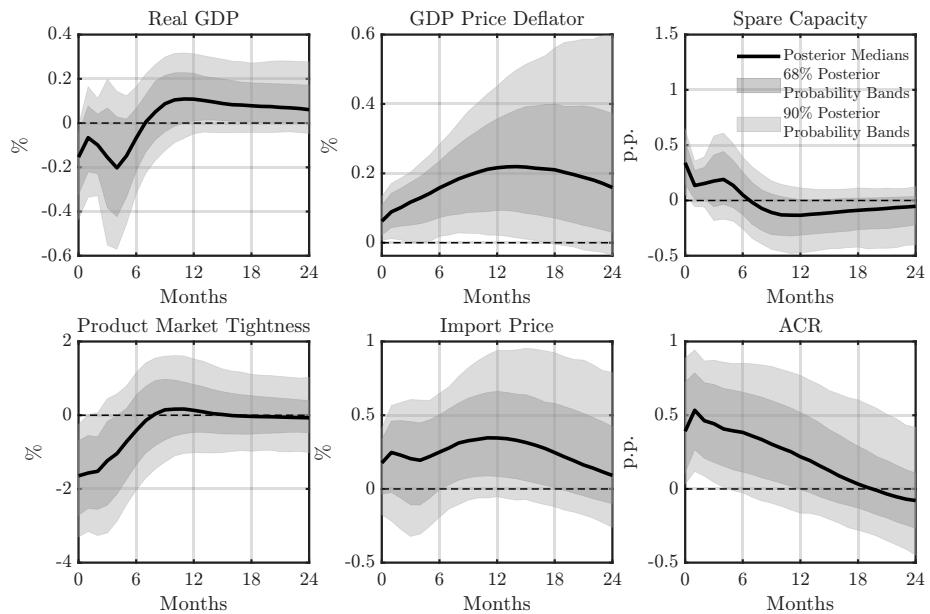


Figure G.42: IRFs to an Adverse Supply Chain Shock: Real GDP and GDP Deflator

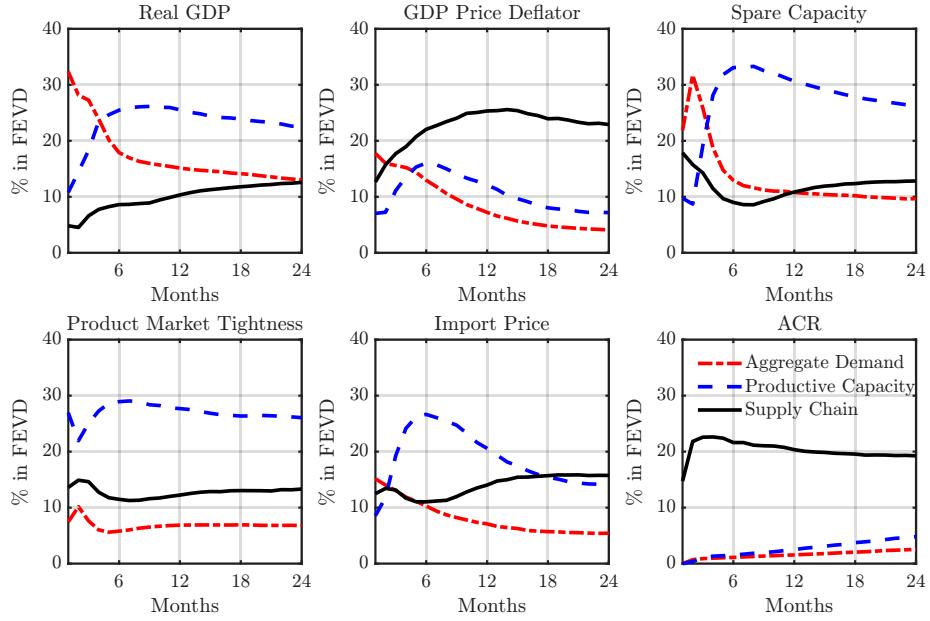


Figure G.43: FEVDs from the SVAR: Real GDP and GDP Deflator

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that real PCE and the PCE price index are replaced by real GDP and the GDP price deflator.

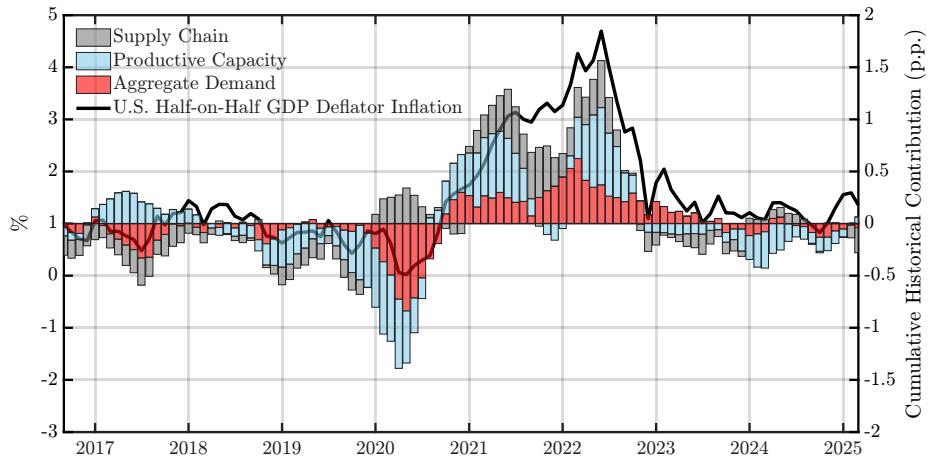


Figure G.44: HD of U.S. HoH GDP Deflator Inflation

Notes. The solid line plots U.S. GDP deflator inflation, measured as the HoH growth rate of the GDP deflator, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks, plotted against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that real PCE and the PCE price index are replaced with real GDP and the GDP deflator.

H. Identification Gains From Supply Chain Proxies

This appendix documents the identification gains —both in the magnitude of price responses to supply chain disturbances and in the precision of the posterior estimates— that arise when our congestion indices are used as measures of global supply chain conditions in the SVAR.

As a starting point, Appendix H.1 estimates the causal effects of supply chain shocks when no proxy is included, providing a natural lower bound against which to evaluate the gains from adding proxies and imposing domain-knowledge-based identification restrictions. Appendix H.2 then includes the ACR index but removes the zero restrictions on its responses to aggregate demand and productive capacity shocks at horizon $k = 1$, showing that much of the improvement comes simply from incorporating our congestion indices.

Appendices H.3 and H.4 show that using the ACT index or a targeted ACR index for major Trans-Pacific ports central to the U.S. goods supply chain delivers results that are quantitatively similar to those from the ACR baseline. By contrast, Appendices H.5-H.7 examine HARPEX, the New York Fed’s GSCPI (Benigno et al., 2022), and the SDI from Smirnyagin and Tsyvinski (2022) and Liu et al. (2024) as alternative proxies, revealing sizable discrepancies that materially alter the estimated macroeconomic effects of supply chain disturbances.

Finally, Appendix H.8 replaces the PCE chain-type price index with a goods-only measure, confirming that the main conclusions remain intact.

H.1. No Proxy

We re-estimate the SVAR without a direct proxy for global supply chain conditions and compare it with the baseline specification that includes ACR under zero restrictions on its responses to aggregate demand and productive capacity shocks at horizon $k = 1$ (Restrictions 1-3). The price IRFs in Figure H.1 clearly illustrate the loss of identification: without a proxy, the median responses of the PCE chain-type price index and the import price index to an adverse supply chain shock are smaller, and the 68% and 90% posterior bands overlap zero across most horizons. By contrast, with ACR and zero restrictions (Figure 9 in the main text), the median responses are larger, the bands are tighter, and they remain above zero for longer horizons, yielding a sharper inflationary signal from supply chain disturbances.

The FEVDs in Figure H.2 convey the same pattern. In the no-proxy case, a larger share of the forecast error variance of PCE prices is attributed to aggregate demand, with only a modest

role for the supply chain shock. In the baseline with ACR and zero restrictions (Figure 10), the supply chain contribution rises at medium horizons, which shows that ACR isolates the relevant supply side component that would otherwise be absorbed by aggregate demand.

The HD of U.S. headline PCE inflation in Figure H.3 reinforces this point. Without a proxy, the 2021-22 inflation surge is attributed mainly to demand and capacity shocks, and the supply chain shock plays a limited role. In the baseline (Figure 11), the supply chain contribution grows visibly larger over the same period, consistent with congestion in maritime logistics and showing that ACR provides cleaner identification of supply-chain-driven price pressures.

Taken together, the no-proxy specification serves as a natural lower bound. Introducing ACR and imposing the $k = 1$ zero restrictions, motivated by industrial practice, increases the magnitude and precision of price responses and shifts the variance and historical contributions toward the supply chain shock, in line with the mechanism emphasized in our framework.

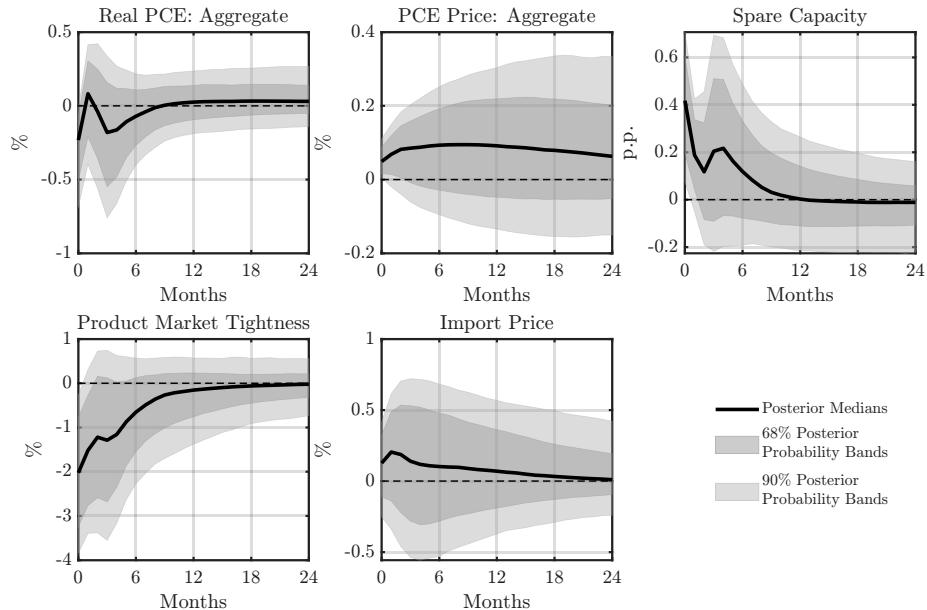


Figure H.1: IRFs to an Adverse Supply Chain Shock: No Proxy

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification described in Section 4, except that the ACR index is omitted.

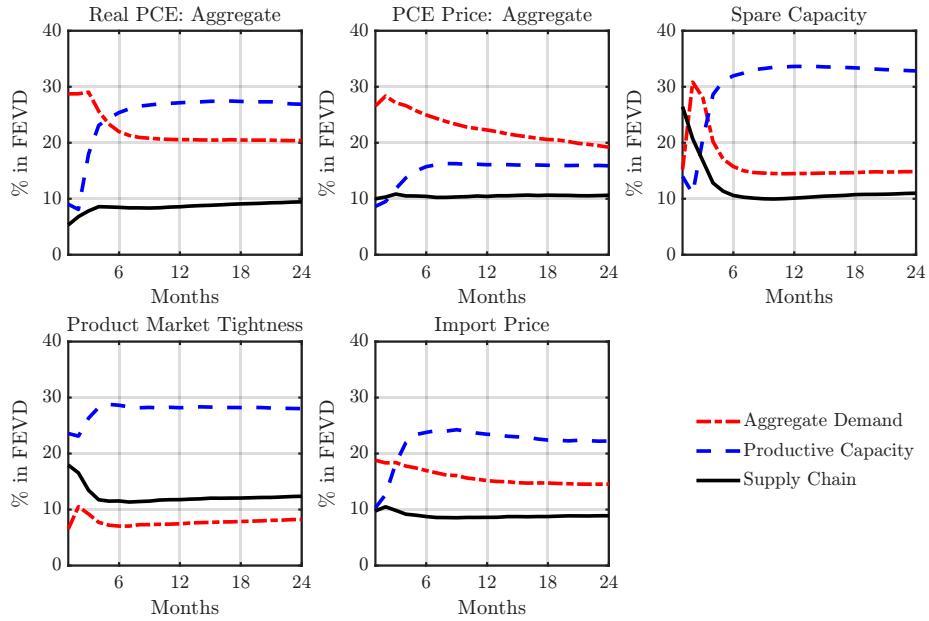


Figure H.2: FEVDs from the SVAR: No Proxy

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the ACR index is omitted.

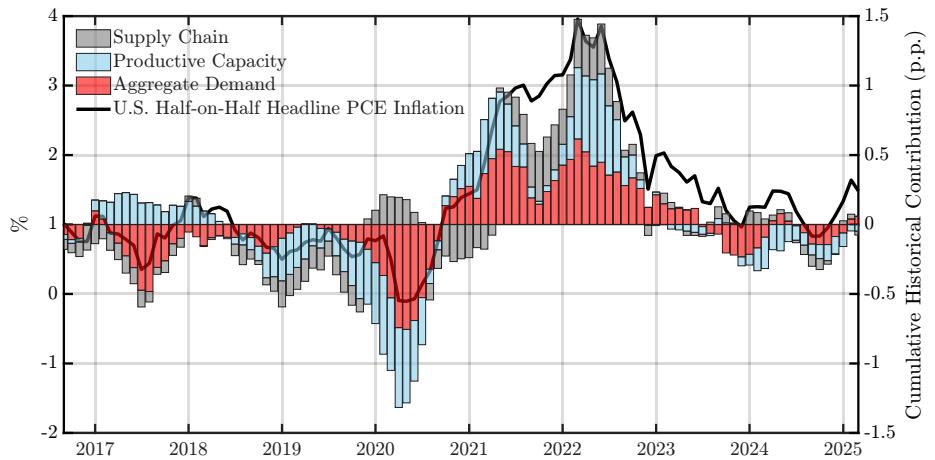


Figure H.3: HD of U.S. HoH Headline PCE Inflation: No Proxy

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR identical to the baseline, except that the ACR index is omitted.

H.2. ACR Without Zero Restrictions

Next, we include the ACR index in the SVAR but remove the zero restrictions on its responses to aggregate demand and productive capacity shocks at horizon $k = 1$. Figure H.4 shows that, relative to the no-proxy case, adding ACR raises the median response of the PCE chain-type price index to an adverse supply chain shock, with narrower 68% and 90% posterior bands that more often lie above zero. However, the responses remain less precisely estimated than in the baseline with ACR and zero restrictions (Figure 9 in the main text), where the bands tighten further and the inflationary signal is clearer.

Figures H.5 and H.6 point to the same conclusion. With the ACR index, the supply chain shock explains a larger share of the forecast error variance of PCE prices and contributes more to the 2021–22 inflation surge than in the no-proxy case. Yet, in both dimensions, its role remains smaller than in the baseline (Figures 10 and 11), where the timing and magnitude more closely align with observed congestion in maritime logistics.

In sum, including ACR alone accounts for most of the identification gains (both in the size of price responses and the precision of estimates), while the zero restrictions provide additional tightening that sharpens the inflationary signal from supply chain shocks.

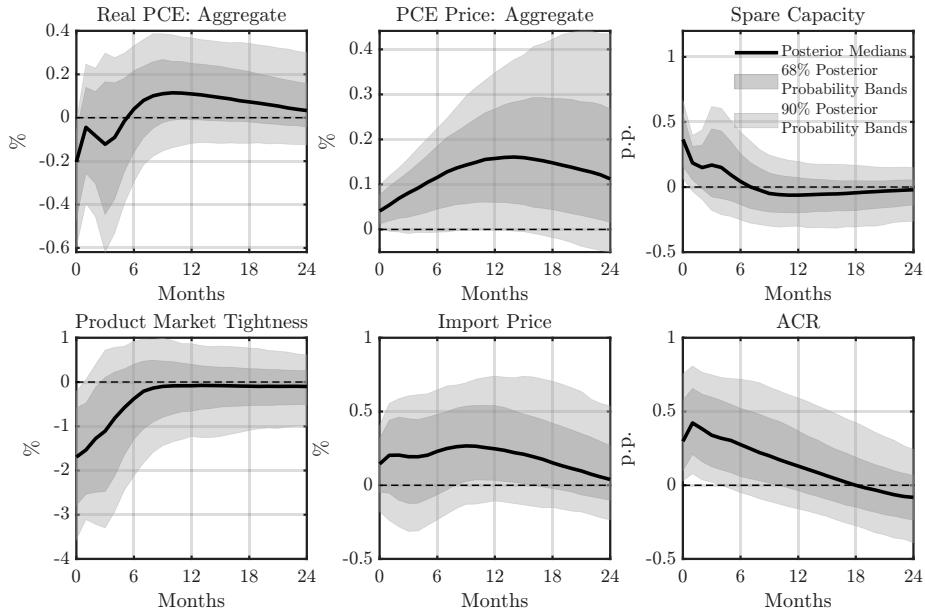


Figure H.4: IRFs to an Adverse Supply Chain Shock: ACR Without Zero Restrictions

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification described in Section 4, except that the zero restrictions on the ACR index are omitted.

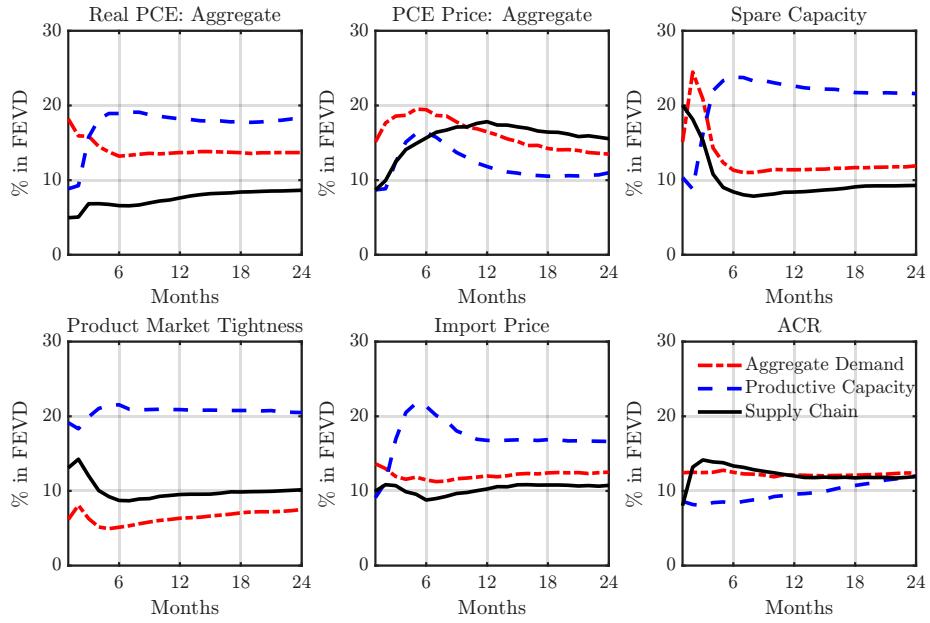


Figure H.5: FEVDs from the SVAR: ACR Without Zero Restrictions

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the zero restrictions on the ACR index are omitted.

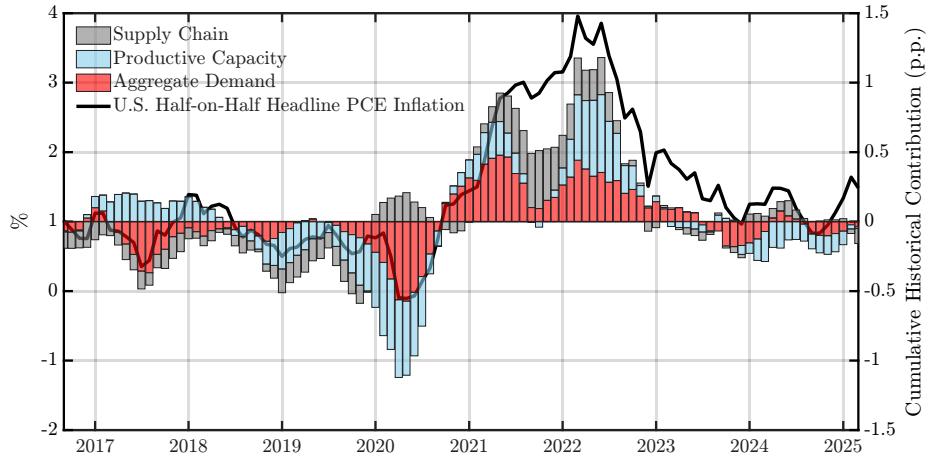


Figure H.6: HD of U.S. HoH Headline PCE Inflation: ACR Without Zero Restrictions

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the zero restrictions on the ACR index are omitted.

H.3. ACT

As introduced in Appendix C.3, we define an alternative port congestion metric, the ACT index. This index measures the average number of hours a container ship waits in a port's anchorage area before docking at a berth, weighted by each port's share of ship visits. Figure H.7 plots the ACT index alongside the ACR index at a monthly frequency. The two series co-move closely and display similar dynamics, except during 2016–2017, when infrastructure upgrades worldwide appear to have focused on the extensive margin (reducing the share of delayed ships) with limited progress on the intensive margin (reducing average delays).¹⁸

Given the high correlation between ACR and ACT, it is unsurprising that, as shown in Figures H.8–H.10, using the ACT index in the causality assessment yields results that are quantitatively similar to those in the ACR baseline.

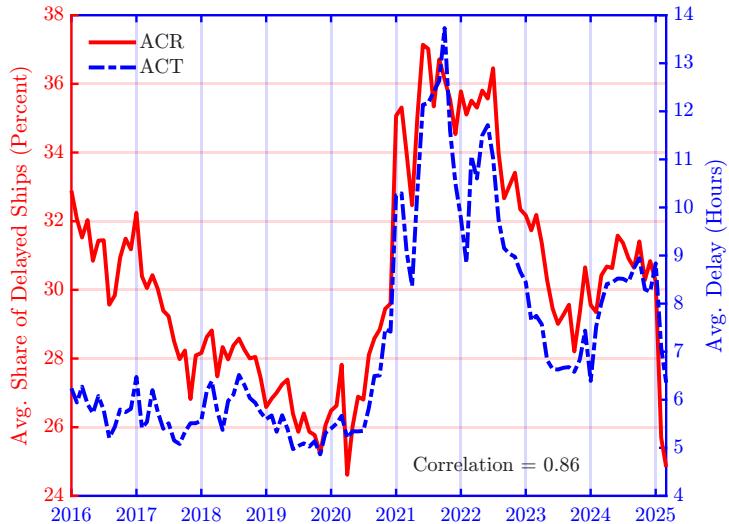


Figure H.7: ACR vs. ACT

Notes. The ACR and the ACT over the sample period January 2016–March 2025 (correlation ≈ 0.86). The ACR is expressed in percentage terms and shown on the left axis, while the ACT is plotted on the right axis in hours. Both indices are seasonally adjusted.

¹⁸In 2016–2017, capacity expansions (new berths, deeper channels, larger cranes) lowered the likelihood of delays at the extensive margin. By contrast, operational reforms needed to shorten average waiting times at the intensive margin were slower to materialize.

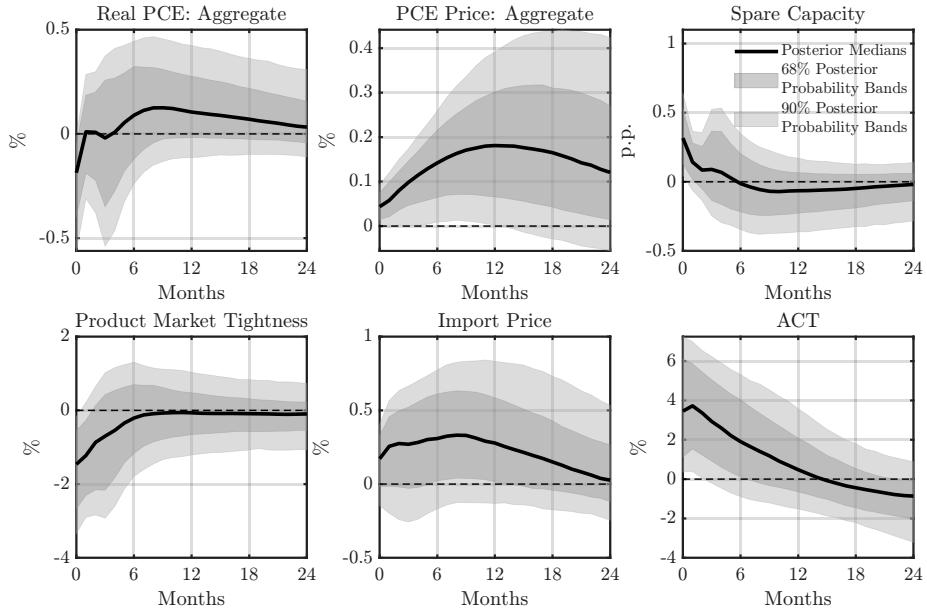


Figure H.8: IRFs to an Adverse Supply Chain Shock: ACT

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification in Section 4, except that the ACT index replaces the ACR index as the proxy for global supply chain conditions.

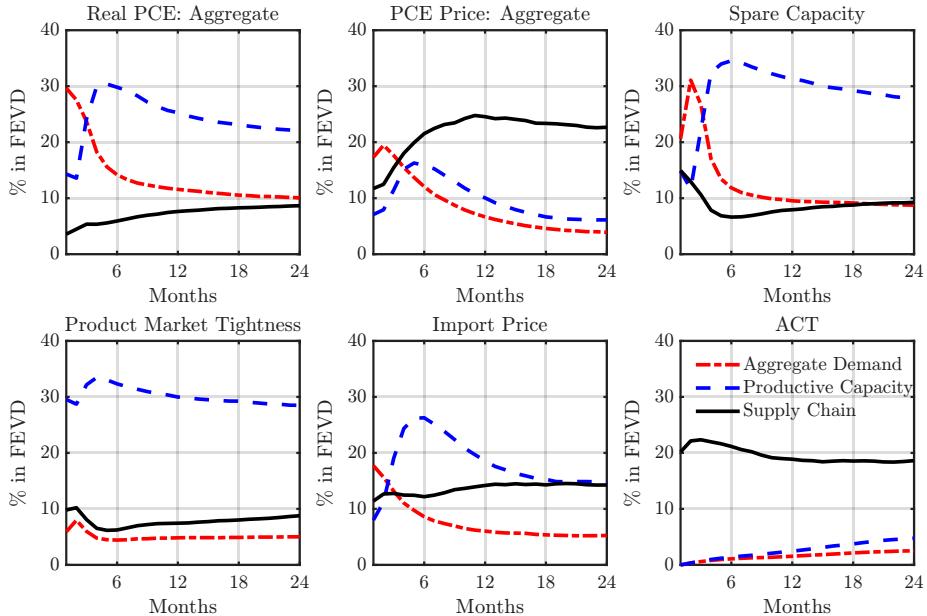


Figure H.9: FEVDs from the SVAR: ACT

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the ACT index replaces the ACR index as the proxy for global supply chain conditions.

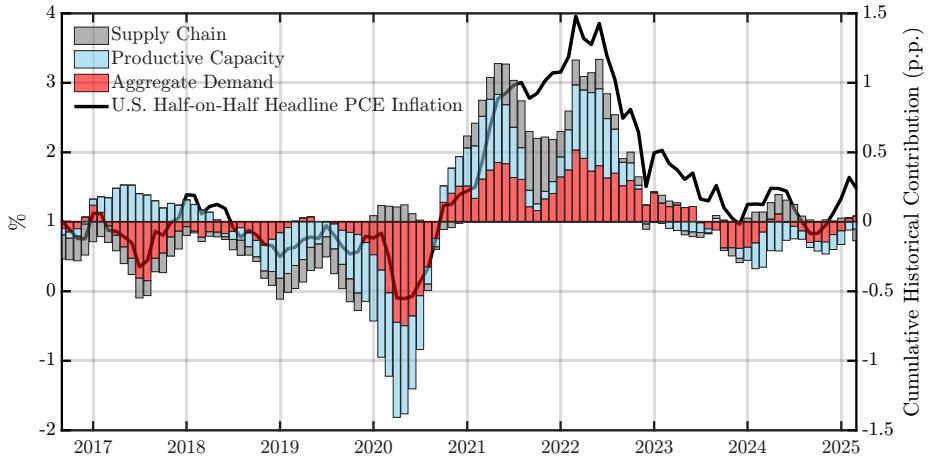


Figure H.10: HD of U.S. HoH Headline PCE Inflation: ACT

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the ACT index replaces the ACR index as the proxy for global supply chain conditions.

H.4. Trans-Pacific ACR

In addition to the global ACR index, we construct a targeted ACR index for the major container ports along the Trans-Pacific route, one of the busiest shipping corridors linking East Asia (primarily China) and the U.S., handling nearly 30 million TEUs annually.¹⁹ We use this index in the SVAR as a robustness check, since congestion along the Trans-Pacific route may exert a disproportionately large influence on the U.S. economy. Shocks identified with this targeted measure may thus reflect supply chain disturbances more specific to the U.S. than those captured by the global ACR index.²⁰

Figure H.11 plots the Trans-Pacific ACR index alongside the global ACR index at a monthly frequency. While the two series track each other closely over the sample, the targeted index displays larger short-term swings, particularly during the pandemic.

The estimated effects are similar to those in the ACR baseline. Figure H.12 shows that a supply chain shock generates stagflationary pressures, raises spare capacity, and lowers product

¹⁹The Trans-Pacific ACR index aggregates congestion at the Port of Busan in South Korea; the Ports of Hanshin and Keihin in Japan; the Ports of Los Angeles and Long Beach in the U.S.; and the Ports of Dalian, Dongguan, Guangzhou Harbor, Hong Kong, Kaohsiung, Ningbo-Zhoushan, Qingdao, Shenzhen, Shanghai, Tianjin, Xiamen, and Yang Shan in China.

²⁰Relatedly, Kilian et al. (2023) develop a monthly index of container trade volumes to and from North America. Although this measure helps identify shocks to domestic demand and to foreign demand for U.S. manufactured goods, it is not suitable for our purpose of identifying supply chain disturbances.

market tightness. The median inflation responses remain sizable and broadly comparable to the baseline, with the lower bounds of the posterior bands staying above zero within a year of the shock. As shown in Figure H.13, supply chain shocks continue to account for the largest share of the forecast error variance of the PCE price index at medium and long horizons, while capacity shocks remain the dominant driver of the other endogenous variables except for the Trans-Pacific ACR. Finally, the HD in Figure H.14 is nearly indistinguishable from that obtained with the global ACR, indicating that the narratives in Section 4 regarding the drivers of U.S. headline PCE inflation are robust, even when focusing specifically on congestion at Trans-Pacific ports.

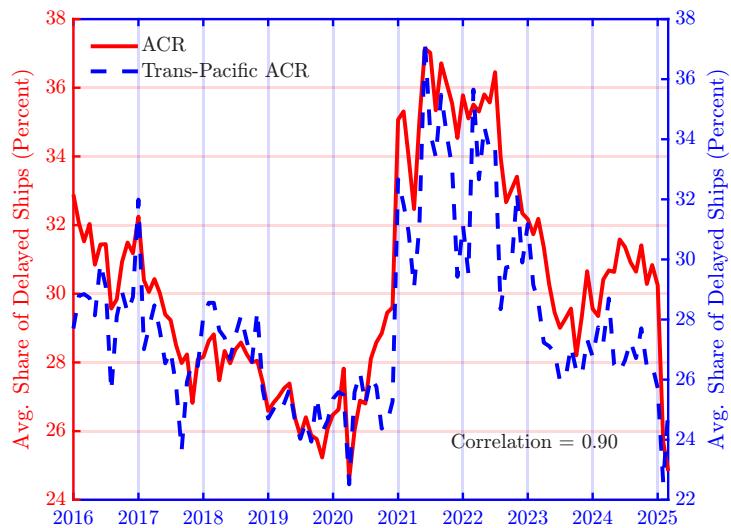


Figure H.11: Global ACR vs. Trans-Pacific ACR

Notes. The global ACR and the targeted ACR for major ports along the Trans-Pacific route over January 2016–March 2025 (correlation ≈ 0.90). Both indices are expressed in percentage terms and seasonally adjusted.

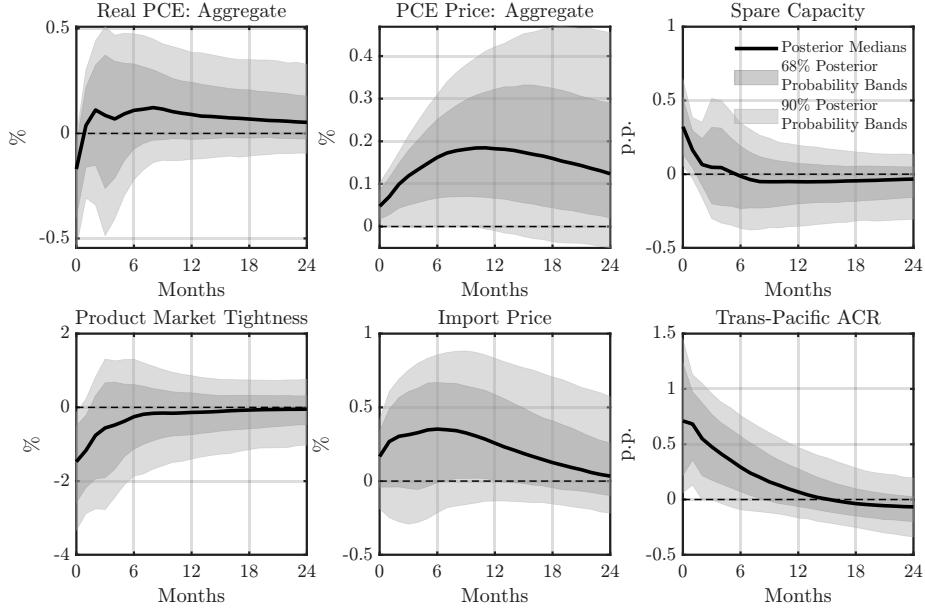


Figure H.12: IRFs to an Adverse Supply Chain Shock: Trans-Pacific ACR

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR described in Section 4, except that the targeted ACR index for major Trans-Pacific ports is used as the proxy for the U.S.-specific supply chain conditions.

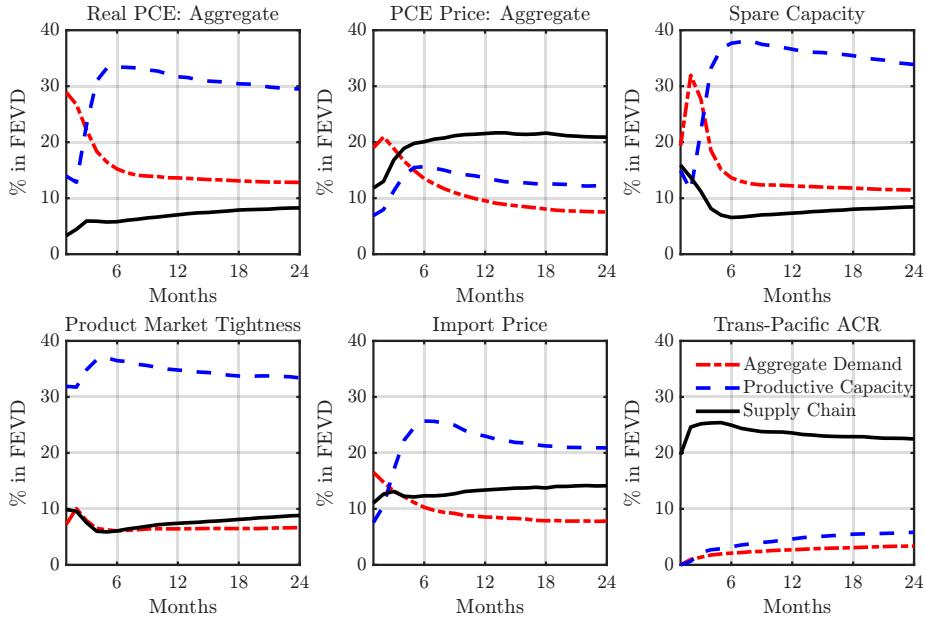


Figure H.13: FEVDs from the SVAR: Trans-Pacific ACR

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable explained by each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR with the same specification as the baseline, except that the targeted ACR index for major Trans-Pacific ports is used as the proxy for the U.S.-specific supply chain conditions.

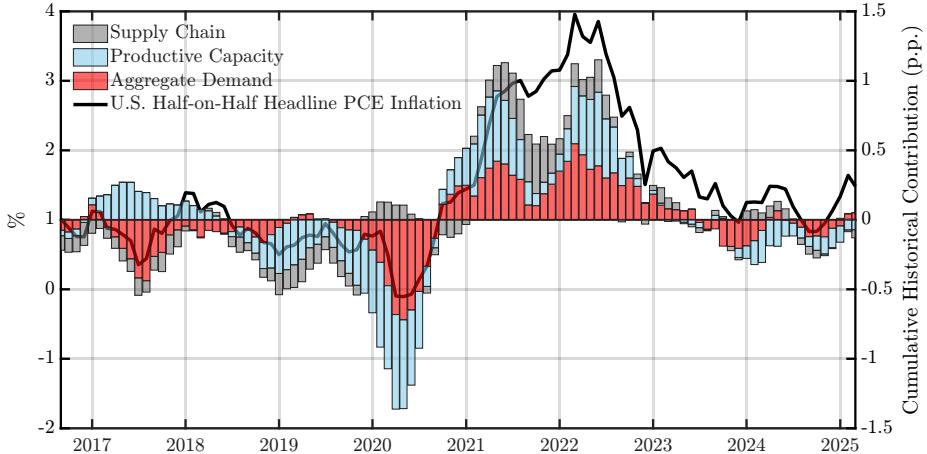


Figure H.14: HD of U.S. HoH Headline PCE Inflation: Trans-Pacific ACR

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the targeted ACR index for major Trans-Pacific ports is used as the proxy for the U.S.-specific supply chain conditions.

H.5. HARPEX

Shipping costs offer a near-real-time read on logistics conditions and are often used to proxy the state of the global supply chain. Within this class, HARPEX (an aggregate of container time-charter rates across vessel sizes) captures the price of shipping capacity rather than congestion itself, and it also enters the New York Fed’s GSCPI (Benigno et al., 2022).²¹ Because charter rates adjust quickly to shifts in demand, fuel costs, and market expectations, we do not impose the zero restrictions on HARPEX’s responses to aggregate demand and productive capacity shocks at horizon $k = 1$. For context, Figure H.15 compares ACR and HARPEX; although the two series co-move, they differ in magnitude and timing, consistent with ACR tracking port congestion and HARPEX reflecting the price of transport services.

Figure H.16 shows that replacing ACR with HARPEX yields a more muted response of the PCE chain-type price index to an adverse supply chain shock: the median effect is small, and the 90% posterior band largely includes zero across horizons. Consistently, the FEVDs in Figure H.17 attribute a larger share of the forecast error variance of PCE prices to aggregate demand at

²¹Finck and Tillmann (2022) also incorporate HARPEX (together with the RWI/ISL container throughput index and the GSCPI) as endogenous variables in SVARs to capture international container shipping and global supply chain conditions. Their identification of a global supply chain shock combines conventional sign restrictions with narrative information (e.g., the 2011 Tōhoku earthquake, the 2021 Suez Canal blockage, and the 2022 Shanghai backlog).

medium horizons, with only a modest portion explained by the supply chain shock, suggesting that HARPEX blends congestion signals with movements driven by demand and other market forces.

The HD in Figure H.18 reinforces this conclusion. During the 2021-22 inflation surge, the contribution from the supply chain shock is present but materially smaller than under the ACR baseline, while demand and capacity account for a greater share. Taken together, the evidence from impulse responses, variance decompositions, and the historical decomposition of U.S. headline PCE inflation indicates that HARPEX primarily reflects broader shipping market dynamics driven by demand and costs, whereas ACR more directly isolates congestion-related pressures.

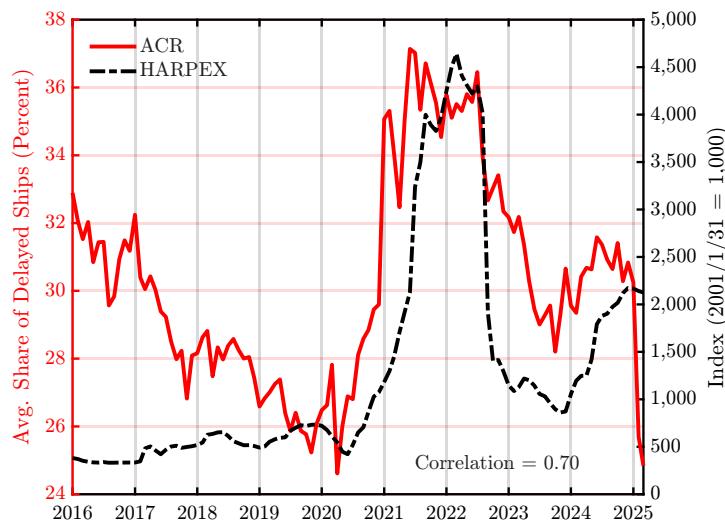


Figure H.15: ACR vs. HARPEX

Notes. The ACR and HARPEX over the sample period January 2016-March 2025 (correlation ≈ 0.70). The ACR is expressed in percentage terms and shown on the left axis, while HARPEX is plotted on the right axis with a baseline value of 1,000 on January 31, 2001. The HARPEX series is published by Harper Peterson and retrieved from the Refinitiv data platform. Both series are seasonally adjusted.

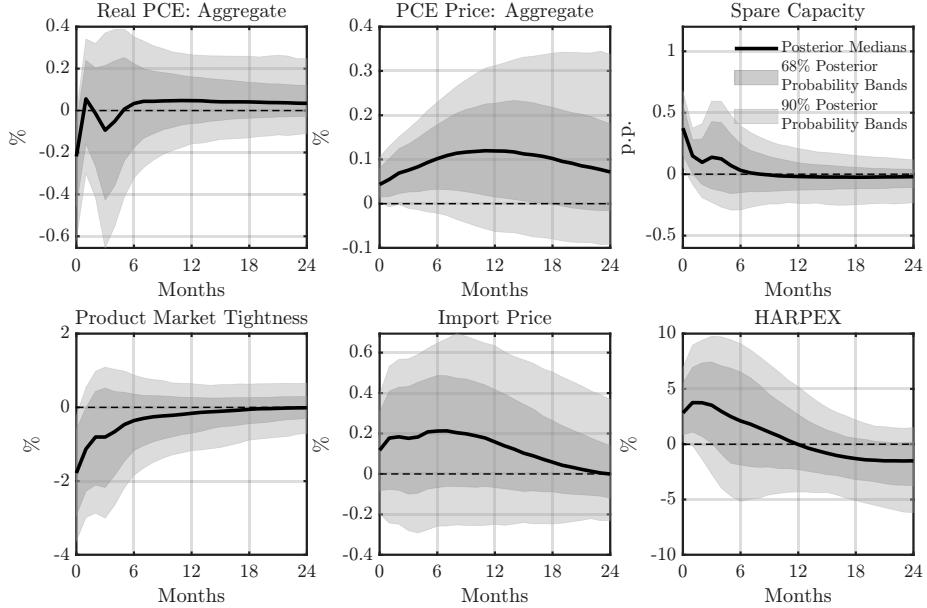


Figure H.16: IRFs to an Adverse Supply Chain Shock: HARPEX

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification described in Section 4, except that HARPEX replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

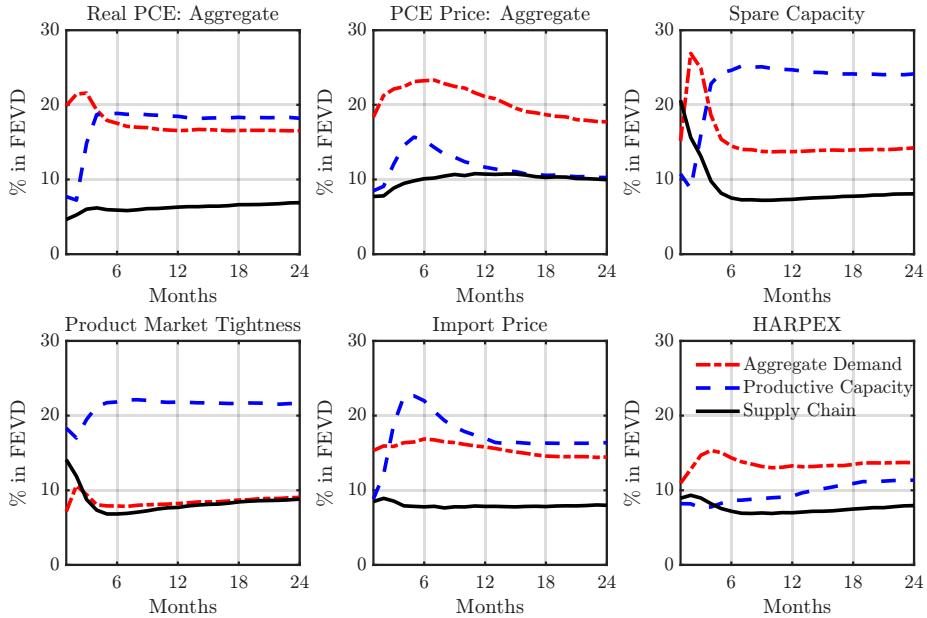


Figure H.17: FEVDs from the SVAR: HARPEX

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that HARPEX replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

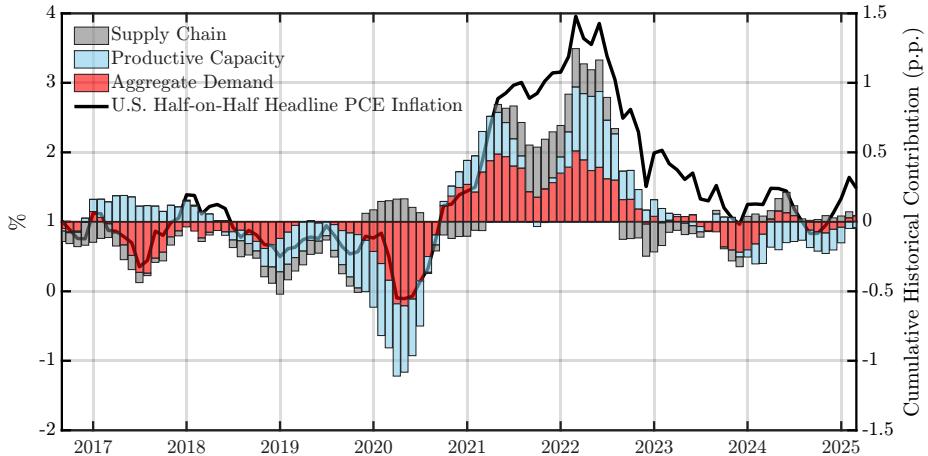


Figure H.18: HD of U.S. HoH Headline PCE Inflation: HARPEX

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that HARPEX replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

H.6. GSCPI

The GSCPI combines cross-border transportation costs and manufacturing PMI subcomponents (e.g., suppliers' delivery times) to infer global supply chain conditions (Benigno et al., 2022). While widely used, these inputs introduce potential measurement errors: transportation costs fluctuate with factors unrelated to supply chain disruptions (such as shifts in fuel prices or shipping demand), PMI indicators reflect subjective managerial assessments that may embed expectations or misperceptions, and longer delivery times can arise from production-side bottlenecks rather than congestion in the logistics network. By contrast, our ACR index is constructed from satellite-based vessel data and directly tracks global port congestion (a first-order manifestation of supply chain disruptions) (Transportation Research Board Executive Committee, 2006).

Figure H.19 compares ACR and the GSCPI over 2016–2025 and shows that, although positively correlated, the two indices diverge in amplitude and timing. During the early pandemic, di Giovanni et al. (2022) interpret the sharp rise in the GSCPI as evidence of China's lockdown and its later decline as reflecting partial reopenings in 2020. The ACR, however, indicates that the lockdown did not generate congestion of a magnitude consistent with severe global disruptions, nor did the reopenings deliver immediate relief at major ports. This contrast foreshadows the

weaker inflation signal obtained when the GSCPI substitutes for ACR.

Figure H.20 shows that the response of the PCE chain-type price index to an adverse supply chain shock is muted: the median effect is small, and the 90% posterior band includes zero throughout, in contrast to the clearer response under ACR with the zero restrictions. Likewise, the FEVDs in Figure H.21 allocate more of the variance in real PCE and in PCE and import prices to aggregate demand at medium horizons, while the supply chain share remains comparatively modest.

The HD in Figure H.22 corroborates this view, showing that the contribution of the supply chain shock to the 2021-22 inflation episode is visible but materially smaller than in the ACR baseline, with demand and capacity accounting for a larger share. Together, these diagnostics point to a diluted imprint of supply chain disturbances when the GSCPI is used in place of ACR.

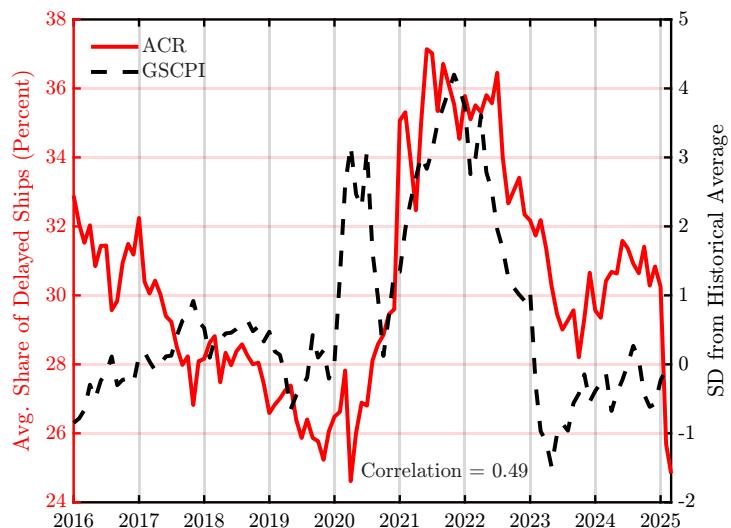


Figure H.19: ACR vs. GSCPI

Notes. The ACR and the New York Fed's GSCPI over January 2016–March 2025 (correlation ≈ 0.49). The ACR is expressed in percentage terms and shown on the left axis, while the GSCPI is plotted on the right axis in standard deviations from its historical average. The GSCPI series is retrieved from the Federal Reserve Bank of New York. Both series are seasonally adjusted.

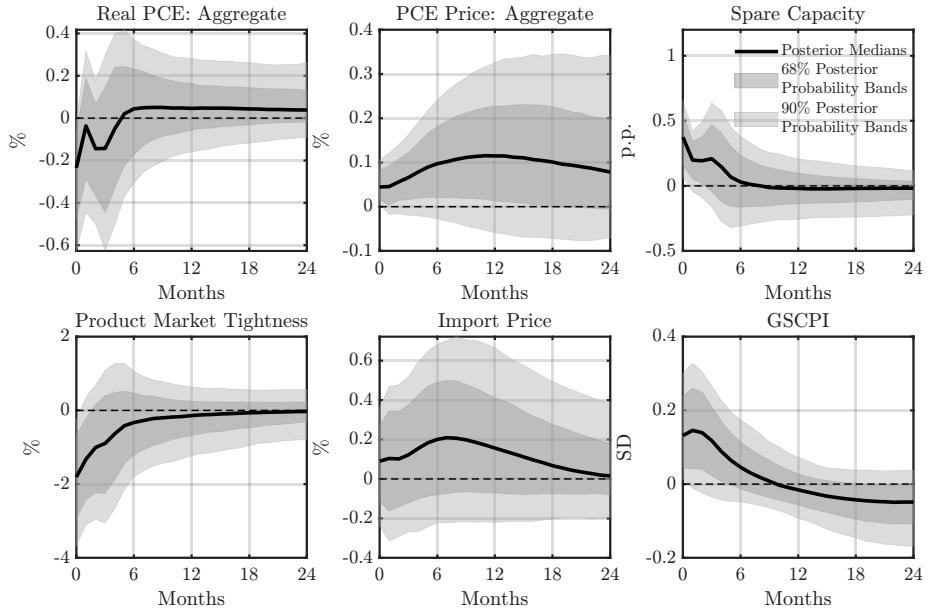


Figure H.20: IRFs to an Adverse Supply Chain Shock: GSCPI

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification described in Section 4, except that the GSCPI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

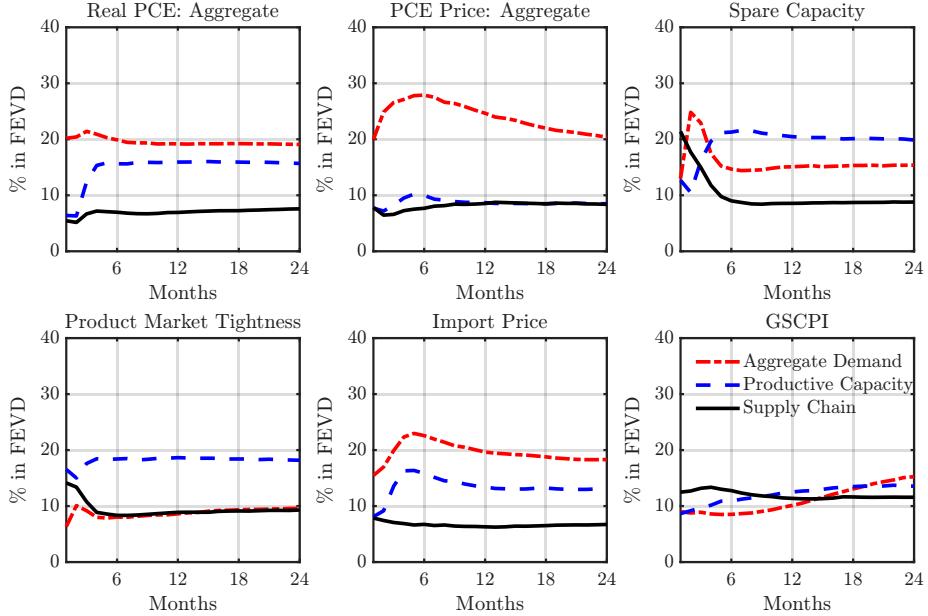


Figure H.21: FEVDs from the SVAR: GSCPI

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the GSCPI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

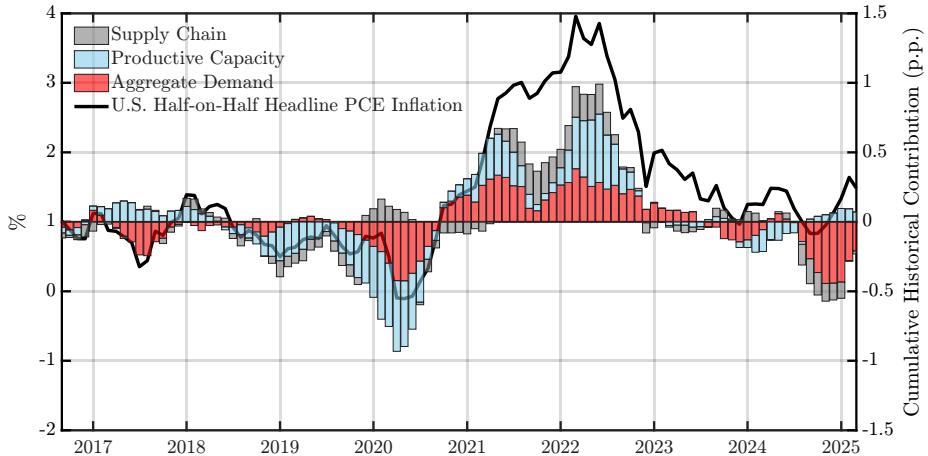


Figure H.22: HD of U.S. HoH Headline PCE Inflation: GSCPI

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the GSCPI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

H.7. Supply Disruptions Index (SDI)

[Smirnyagin and Tsyvinski \(2022\)](#) and [Liu et al. \(2024\)](#) construct the SDI using the S&P Global Panjiva dataset of U.S. seaborne import records. The index tracks regular consignee-shipper relationships on a quarterly basis and records a disruption when an otherwise active pair becomes inactive for one quarter before resuming. From 2016 to 2025, the SDI has remained relatively stable before the pandemic, rose in early 2020, showed renewed elevation during the 2021-22 inflation episode, and then eased thereafter (Figure [H.23](#)).

Turning to the SVAR evidence, Figure [H.24](#) shows that when SDI replaces ACR, the responses of the PCE chain-type price index (and likewise the import price index) to an adverse supply chain shock are less precisely estimated: median effects are small, and the 90% posterior band often includes zero across horizons.

Consistently, the FEVDs in Figure [H.25](#) assign only a small share of fluctuations in PCE and import prices to supply chain shocks at medium horizons. The HD in Figure [H.26](#) likewise shows that supply chain shocks made only a limited contribution to the 2021-22 inflation episode, with demand and capacity playing larger roles. Echoing the results with the GSCPI, substituting SDI for ACR weakens the identification of the inflationary effects of supply chain disturbances.

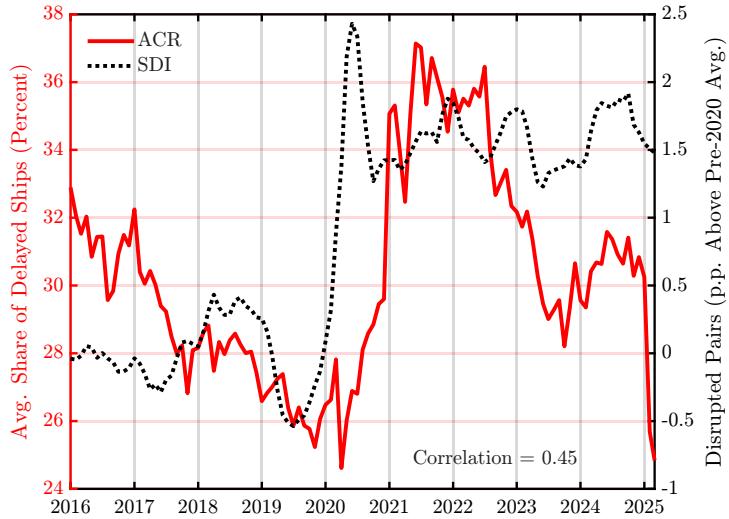


Figure H.23: ACR vs. SDI

Notes. The ACR and the SDI over the sample period January 2016–March 2025 (correlation ≈ 0.45). The ACR is expressed as a percentage and shown on the left axis, while the SDI is plotted on the right axis in percentage points. The SDI series is retrieved from <https://www.disruptions.supply> (accessed August 3, 2025). Both series are seasonally adjusted.

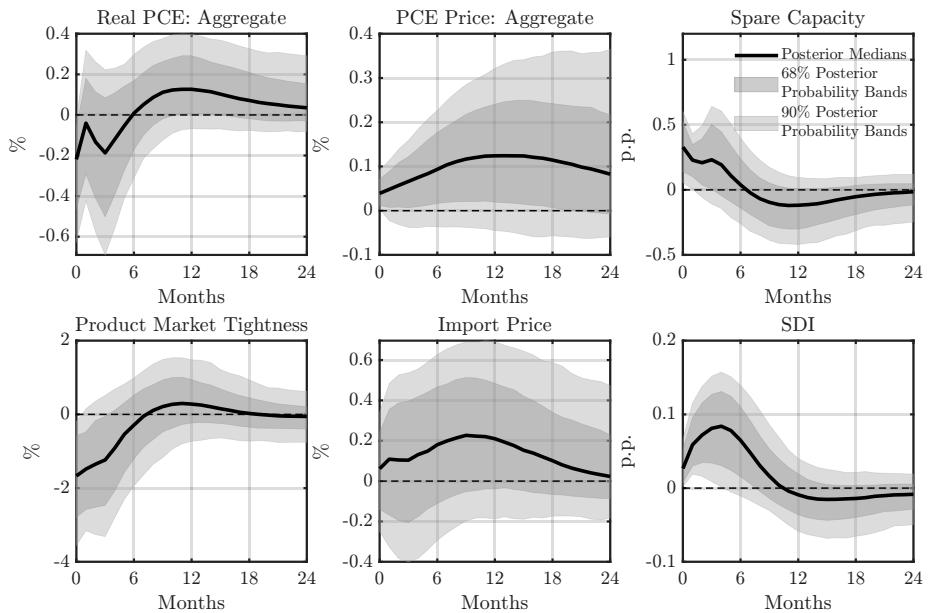


Figure H.24: IRFs to an Adverse Supply Chain Shock: SDI

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification in Section 4, except that the SDI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

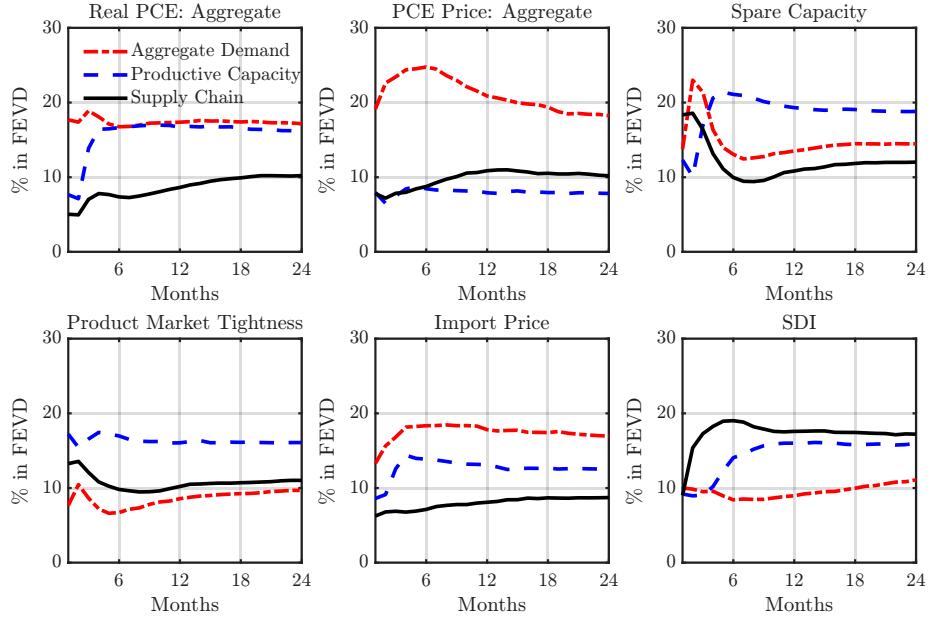


Figure H.25: FEVDs from the SVAR: SDI

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the SDI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

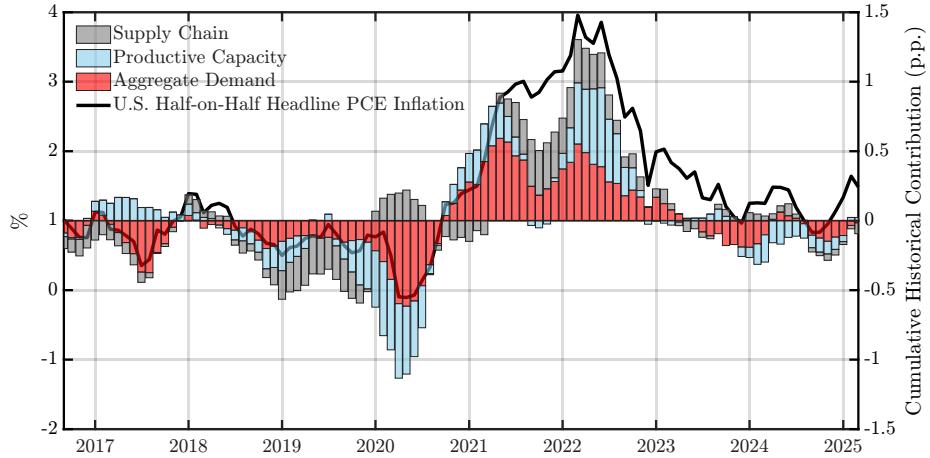


Figure H.26: HD of U.S. HoH Headline PCE Inflation: SDI

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the SDI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

H.8. Goods Price Responses Across Proxies

We replicate the cross-proxy exercise in Section 4.4 using the PCE goods price index instead of the aggregate PCE price index. Figure H.27 reports impulse responses to a one-standard-deviation adverse supply chain shock across eight specifications: no proxy; ACR; ACR with zero restrictions on its responses to demand and capacity shocks at horizon $k = 1$; ACT with the same restrictions; Trans-Pacific ACR with the same restrictions; HARPEX; GSCPI; and SDI. When ACR, ACT, or the Trans-Pacific ACR is combined with these zero restrictions, the median responses are larger and more precisely estimated —posterior bands are tighter, with the 68% interval above zero at nearly all horizons and the 90% interval largely above zero—clearly indicating the inflationary nature of supply chain disturbances and consistent with our theoretical predictions. In contrast, omitting a proxy or using HARPEX, GSCPI, or SDI yields smaller responses with wider bands that generally include zero.

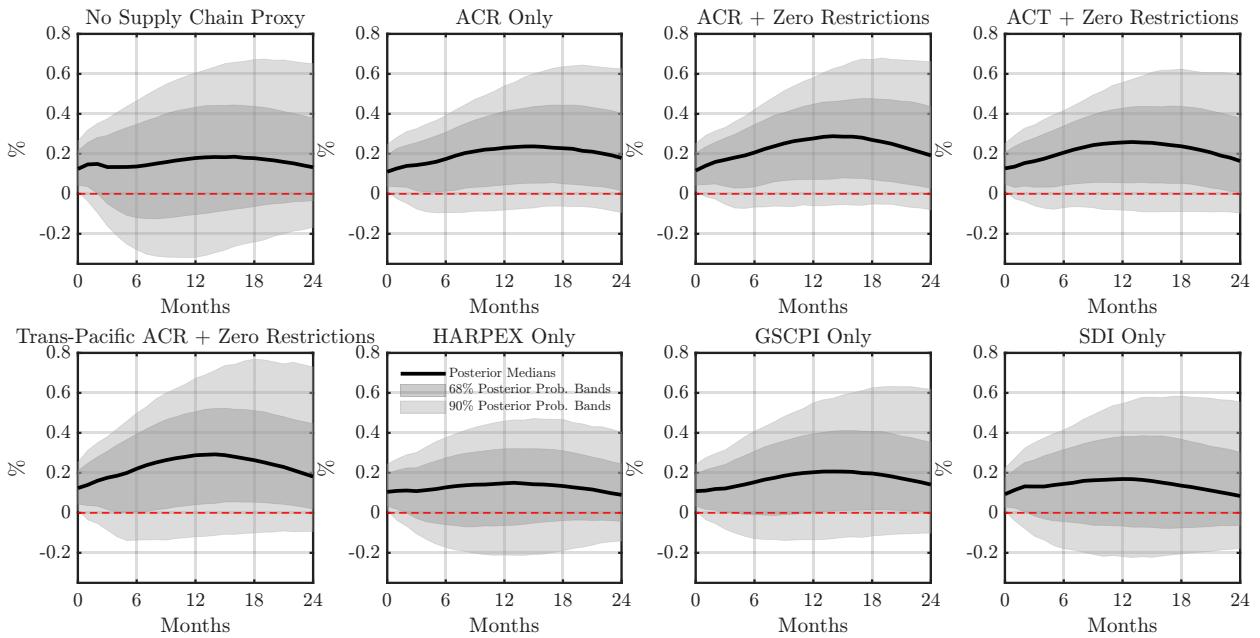


Figure H.27: Goods Price Responses to an Adverse Supply Chain Shock Across Proxies

Notes. Posterior-median impulse responses of the U.S. PCE goods price index, with 68% and 90% posterior probability bands, to a one-standard-deviation adverse supply chain shock across eight SVAR specifications: (i) no supply chain proxy; (ii) ACR; (iii) ACR with zero restrictions at horizon $k = 1$ on its responses to aggregate demand and productive capacity shocks; (iv) ACT with the same zero restrictions; (v) Trans-Pacific ACR with the same zero restrictions; (vi) HARPEX; (vii) GSCPI; and (viii) SDI. Aside from the choice of proxy, all eight specifications follow the baseline model in Section 4, except that real PCE and the aggregate PCE price index are replaced by goods PCE and the PCE goods price index. The sign restrictions in Restrictions 1-3 are imposed to identify the adverse supply chain shock in all specifications except (i), where the positive restriction on the ACR response is omitted.

The FEVD results in Figure H.28 show that the congestion indices, together with the zero restrictions motivated by our domain knowledge, assign a larger share of the forecast error variance of PCE goods prices to supply chain disturbances at medium horizons. In contrast, with alternative proxies, the attribution shifts toward aggregate demand. These patterns mirror the aggregate-price results and highlight the identification gains from combining congestion-based proxies with zero restrictions.

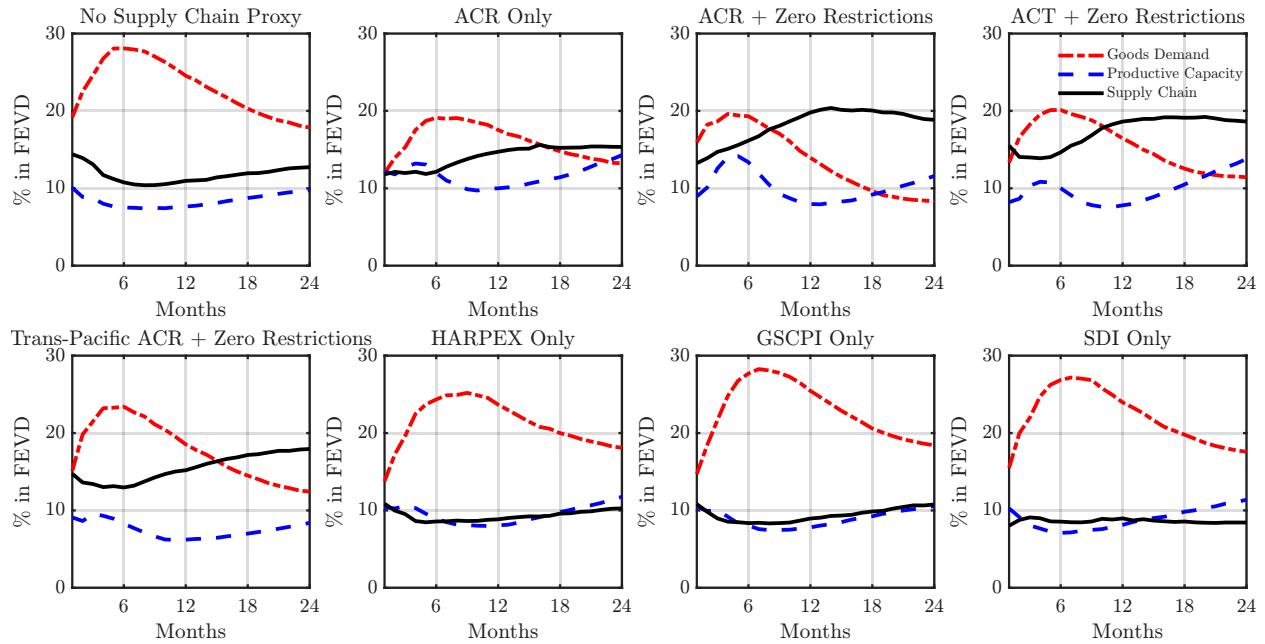


Figure H.28: FEVD of the U.S. PCE Goods Price Index Across Proxies

Notes. Posterior-median shares of the forecast error variance of the U.S. PCE goods price index attributable to goods demand, productive capacity, and supply chain shocks across horizons for each of the eight specifications discussed above.

References for Appendices

Allen, T. (2014). Information Frictions in Trade. *Econometrica*, 82:2041–2083.

Arias, J. E., Rubio-Ramírez, J. F., and Waggoner, D. F. (2018). Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications. *Econometrica*, 86:685–720.

Bai, X., Ma, Z., Hou, Y., Li, Y., and Yang, D. (2023). A Data-Driven Iterative Multi-Attribute Clustering Algorithm and Its Application in Port Congestion Estimation. *IEEE Transactions on Intelligent Transportation Systems*, 24:12026–12037.

Baqae, D. and Farhi, E. (2019). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem. *Econometrica*, 87(4):1155–1203.

Baqae, D. and Farhi, E. (2022). Supply and Demand in Disaggregated Keynesian Economies With an Application to the COVID-19 Crisis. *American Economic Review*, 112:1397–1436.

Barnichon, R. and Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *Review of Economics and Statistics*, 101(3):522–530.

Bauer, M. D. and Swanson, E. T. (2023). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. *NBER Macroeconomics Annual*, 37:87–155.

Benguria, F. (2021). The Matching and Sorting of Exporting and Importing Firms: Theory and Evidence. *Journal of International Economics*, 131:103430.

Benigno, G., di Giovanni, J., Groen, J. J., and Noble, A. I. (2022). The GSCPI: A New Barometer of Global Supply Chain Pressures. Staff Report 1017, Federal Reserve Bank of New York.

Bils, M., Chang, Y., and Kim, S.-B. (2011). Worker Heterogeneity and Endogenous Separations in a Matching Model of Unemployment Fluctuations. *American Economic Journal: Macroeconomics*, 3:128–154.

Birant, D. and Kut, A. (2007). ST-DBSCAN: An Algorithm for Clustering Spatial–Temporal Data. *Data & Knowledge Engineering*, 60:208–221.

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2020). Geography, Transportation, and Endogenous Trade Costs. *Econometrica*, 88(2):657–691.

Chaney, T. (2014). The Network Structure of International Trade. *American Economic Review*, 104(11):3600–3634.

Chow, G. C. and Lin, A.-l. (1971). Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series. *Review of Economics and Statistics*, 53:372–375.

di Giovanni, J., Sebnem Kalemli-Özcan, Silva, A., and Yildirim, M. A. (2022). Global Supply Chain Pressures, International Trade, and Inflation. Working Paper 30240, National Bureau of Economic Research.

Du, Y., Chen, Q., Lam, J. S. L., Xu, Y., and Cao, J. X. (2015). Modeling the Impacts of Tides and the Virtual Arrival Policy in Berth Allocation. *Transportation Science*, 49(4):939–956.

Eaton, J. and Kortum, S. (2002). Technology, Geography, and Trade. *Econometrica*, 70:1741–1779.

Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD-96)*, pages 226–231. AAAI Press.

Fernández-Villaverde, J., Mandelman, F., Yu, Y., and Zanetti, F. (2024). Search Complementarities, Aggregate Fluctuations, and Fiscal Policy. *Review of Economic Studies*, page rdae053.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., Sargent, T. J., and Watson, M. W. (2007). ABCs (and Ds) of Understanding VARs. *American Economic Review*, 97(3):1021–1026.

Finck, D. and Tillmann, P. (2022). The Macroeconomic Effects of Global Supply Chain Disruptions. Discussion Paper 14/2022, Bank of Finland Institute for Emerging Economies.

Fujita, S. and Ramey, G. (2012). Exogenous Versus Endogenous Separation. *American Economic Journal: Macroeconomics*, 4:68–93.

Giacomini, R. and Kitagawa, T. (2021). Robust Bayesian Inference for Set-Identified Models. *Econometrica*, 89:1519–1556.

Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

Kasahara, H. and Lapham, B. (2013). Productivity and the Decision to Import and Export: Theory and Evidence. *Journal of International Economics*, 89:297–316.

Kilian, L., Nomikos, N., and Zhou, X. (2023). Container Trade and the U.S. Recovery. *International Journal of Central Banking*, 19(1):417–450.

Krolikowski, P. M. and McCallum, A. H. (2021). Goods-Market Frictions and International Trade. *Journal of International Economics*, 129:103411.

Lenoir, C., Martin, J., and Mejean, I. (2022). Search Frictions in International Goods Markets. *Journal of the European Economic Association*, 21(1):326–366.

Li, C., Qi, X., and Song, D. (2016). Real-Time Schedule Recovery in Liner Shipping Service With Regular Uncertainties and Disruption Events. *Transportation Research Part B: Methodological*, 93:762–788.

Li, Y., Xu, S., Xia, S., Ma, Z., and Bai, X. (2025). Effective Global Maritime Supply: Determinants, Dynamics, and Economic Impacts. Working paper, Tsinghua University.

Liu, E., Smirnyagin, V., and Tsyvinski, A. (2024). Supply Chain Disruptions and Supplier Capital in U.S. Firms. Working paper, SSRN.

Melitz, J. and Toubal, F. (2014). Native Language, Spoken Language, Translation and Trade. *Journal of International Economics*, 93:351–363.

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, 71:1695–1725.

Menzio, G. and Shi, S. (2011). Efficient Search on the Job and the Business Cycle. *Journal of Political Economy*, 119:468–510.

Naudé, W. and Matthee, M. (2011). The Impact of Transport Costs on New Venture Internationalisation. *Journal of International Entrepreneurship*, 9:62–89.

Newey, W. K. and West, K. D. (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–708.

Notteboom, T. E. (2006). The Time Factor in Liner Shipping Services. *Maritime Economics & Logistics*, 8:19–39.

Rodrigue, J.-P. (2020). *The Geography of Transport Systems*. Routledge, 5 edition.

Smets, F. and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97(3):586–606.

Smirnyagin, V. and Tsyvinski, A. (2022). Macroeconomic and Asset Pricing Effects of Supply Chain Disasters. Working Paper 30503, National Bureau of Economic Research.

Transportation Research Board Executive Committee (2006). *Critical Issues in Transportation*. The National Academies Press.

Wong, W. F. (2022). The Round Trip Effect: Endogenous Transport Costs and International Trade. *American Economic Journal: Applied Economics*, 14:127–166.