American Economic Review 2026, 116(11): 4243–4281
https://doi.org/10.1257/aer.20250081

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

By Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti*

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 US disinflation in 2020, while the inflation surge in 2021 was driven mainly by supply chain shocks. (*JEL* C43, E31, F14, F23, L14, L92)

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. 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.

*Bai: Tsinghua University (email: xiwenbai@mail.tsinghua.edu.cn); Fernández-Villaverde: University of Pennsylvania (email: jesusfv@econ.upenn.edu); Li: Peking University HSBC Business School and UIBE (email: yiliang_li@phbs.pku.edu.cn); Zanetti: University of Oxford (email: francesco.zanetti@economics.ox.ac.uk). Authors are listed alphabetically in this paper's byline. Yuriy Gorodnichenko was the coeditor for this article. We thank the coeditor 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, Zhuhan Jin, Callum Jones, Şebnem Kalemli-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, Katheryn 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. Zanetti gratefully acknowledges financial support from the British Academy Grant SRG2526\261943. Xiwen acknowledges financial support from the National Natural Science Foundation of China (Projects 72371143 and 72188101) and Tsinghua University Initiative Scientific Research Program.

[†]Go to https://doi.org/10.1257/aer.20250081 to visit the article page for additional materials and author disclosure statement(s).

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 (Transportation Research Board Executive Committee 2006), a view later reinforced by Fan, Wilson, and Dahl (2012) and Brancaccio, Kalouptsidi, and Papageorgiou (2024), who documented its impact on the efficiency and reliability of global supply chains.

Container shipments are central to global trade, with roughly 60 percent 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 percent to 37 percent, while average delay durations rose from 6 to 14 hours. Given that almost 80 percent of global trade is shipped indirectly, with an average of five port stops (Ganapati, Wong, and Ziv 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 the condition that the ACR index remains unresponsive to aggregate demand or productive capacity shocks during the

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

first month postimpact, 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, 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 US trading partners (Mexico, Canada, China, Germany, and Japan), which together account for more than half of US goods imports. At the same time, product market tightness falls sharply, proxied by the imbalance between US manufacturers' new orders and the US dollar value of the import-weighted average spare capacity of the same trading partners. This shock also triggers a pronounced decline in US real personal consumption expenditures (PCE) and a persistent increase in the PCE price index, consistent with recent empirical evidence (Khalil 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 US 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, Lindé, and Trabandt 2023), and capacity constraints (Comín, Johnson, and Jones 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, Kalouptsidi, and Papageorgiou 2020; Dunn and Leibovici 2023), infrastructure investment (Fuchs and Wong 2022; Brancaccio, Kalouptsidi, and

Papageorgiou 2024), asset prices (Smirnyagin and Tsyvinski 2022), working capital (Antràs 2023; Kim and Shin 2023), inflation expectations (Acharya, Crosignani, Eisert, and Eufinger 2023; Binetti, Nuzzi, and Stantcheva 2024), the design of new taxes and pricing rules to offset distortionary effects on the transportation network (Brancaccio, Kalouptsidi, Papageorgiou, and Rosaia 2023), the interlinks between oil shocks and trading externalities in the supply chain (Bai and Li 2022; Li, Bai, Wang, and Ma 2022), and the economic effects of supply chain disruptions during the COVID-19 pandemic (Finck and Tillmann 2022; Ascari, Bonam, and Smadu 2024; Finck, Klein, and Tillmann 2024).

The remainder of the paper is organized as follows. Section I constructs the ACR index measuring the state of the global supply chain. Section II develops our theoretical model and the identification restrictions for structural shocks. Section III presents the estimation results from a Bayesian SVAR. Section IV concludes. A Supplemental Appendix provides further details, and our data are available on ICPSR (Bai, Fernández-Villaverde, Li, and Zanetti 2026a).

I. 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.

A. *Containerized Seaborne Trade*: *Some Basic Facts*

Containerized seaborne trade plays a central role in the global supply chain, accounting for about 46 percent of all international trade.[2] In the United States, 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 percent of US trade by weight and roughly 30 percent 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.

As Brancaccio, Kalouptsidi, and Papageorgiou (2020, 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.

[2] See Notteboom, Pallis, and Rodrigue (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).

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 two to three days at several major ports and heavy financial losses.[3] Since nearly 80 percent of world trade is shipped indirectly and the average shipment stops at five intermediate ports before reaching its final destination (Ganapati, Wong, and Ziv 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 ultralarge 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 240,000 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 219,000 tons, about 10 percent 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 restrictions 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 transpacific corridor between East Asia and North America) and increased ship calls to meet demand (see Supplemental Appendix A and Bai et al. 2026 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.

[3] Buyers 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.

[4] For perspective on shipments to the United States, it takes about 28 days to move a container from Shanghai to Los Angeles along the transpacific 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 two to three days at each intermediate port, the total delay amounts to 8 to 12 days, even before accounting for congestion at the origin and destination ports.

[5] "Review of Maritime Transport 2022: Navigating Stormy Waters," UN Trade and Development, accessed December 29, 2023, https://unctad.org/rmt2022.

B. *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, Moxnes, Ulltveit-Moe, and Zi forthcoming). Each record includes the IMO number, time stamp, 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.

C. *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, Kalouptsidi, and Papageorgiou 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, Cerdeiro, and Liu 2022). An anchorage is a location within a port where ships lower their anchors, and 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.

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 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

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

[7] More than 99 percent 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.

[8] Nautical 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, Shaw, Peng, and Fang 2020). Inspecting satellite images can help identify berths with fixed locations and boundaries—for example, Supplemental 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.

Panel A. Headings at a berth

Panel B. Headings at an anchorage

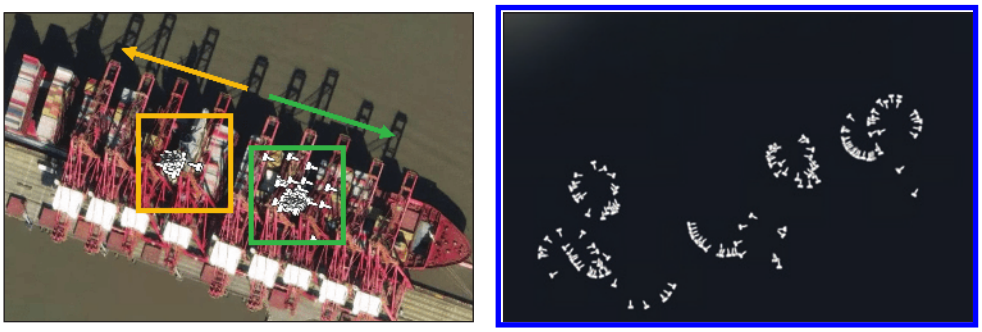


FIGURE 1. INFORMATION ON HEADINGS: TWO EXAMPLES

Note: Imagery © 2020 Planet Labs

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 are random, with some forming a ring shape, consistent with anchorage behavior. Supplemental 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 nonspatial attributes for iterative clustering, provides a versatile tool for classifying clusters with specific labels in diverse contexts (e.g., disease hot spots, 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 Supplemental Appendix C.2.

D. *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

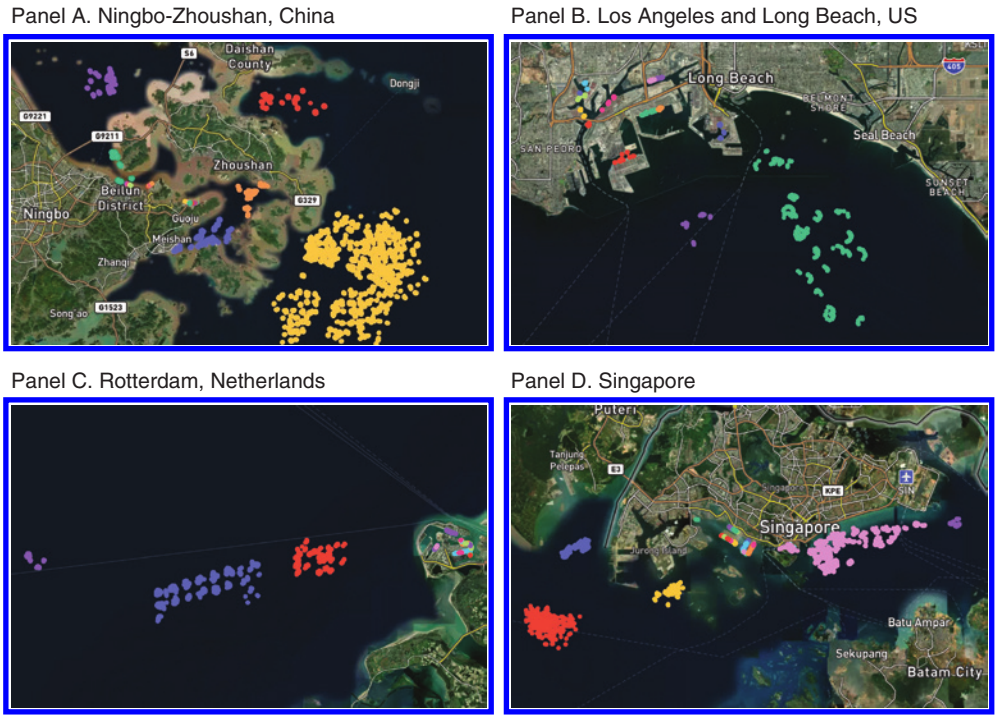


FIGURE 2. IDENTIFICATION OF ANCHORAGE AND BERTH AREAS OF A PORT USING MACHINE LEARNING

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.

Soource: Imagery © 2020 Planet Labs

docking at a berth.[9] These top 50 ports represent 75.6 percent 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:

$$Congestion_{p,t} \equiv \frac{Delayed_{p,t}}{Delayed_{p,t} + Undelayed_{p,t}}, \quad \forall p \in \mathcal{P}, \tag{1}$$

where $Delayed_{p,t}$ and $Undelayed_{p,t}$ 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, Kalouptsidi, and Papageorgiou 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 was 186 hours,

[9] A 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.

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

[11] A 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.

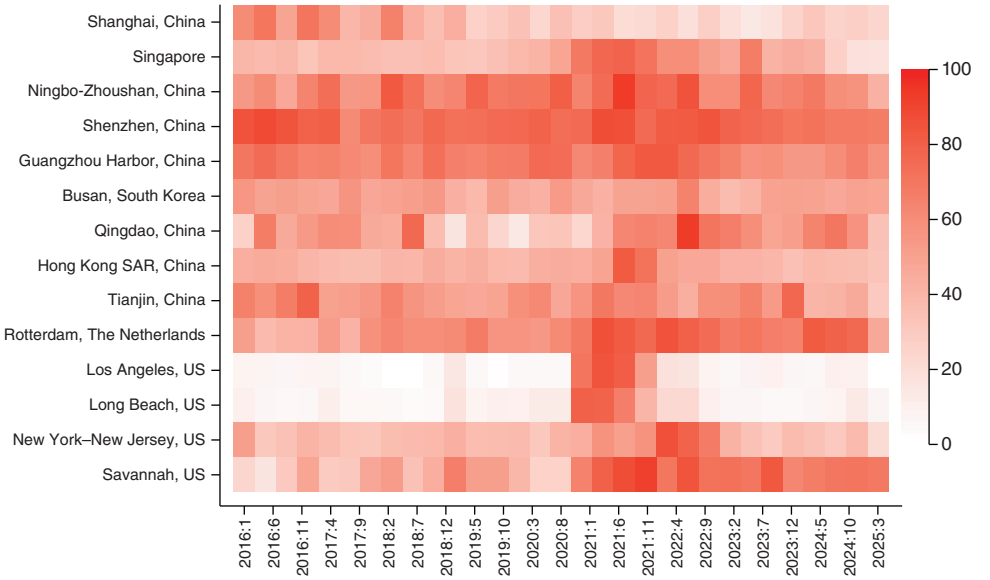


FIGURE 3. CONGESTION RATES FOR THE MAJOR CONTAINER PORTS WORLDWIDE

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.

while in Shanghai it was virtually zero. Moreover, the rate measure 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 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 percent of US 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 percent 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 percent of inbound ships at the Port of Los Angeles were unable to dock immediately upon arrival in late 2020.[13]

[12] Nevertheless, Supplemental Appendix H.3 also constructs an index of Average Congestion Time (ACT). Consistent with Brancaccio, Kalouptsidi, and Papageorgiou (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.

[13] This 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 percent in August to 26 percent in December 2020.

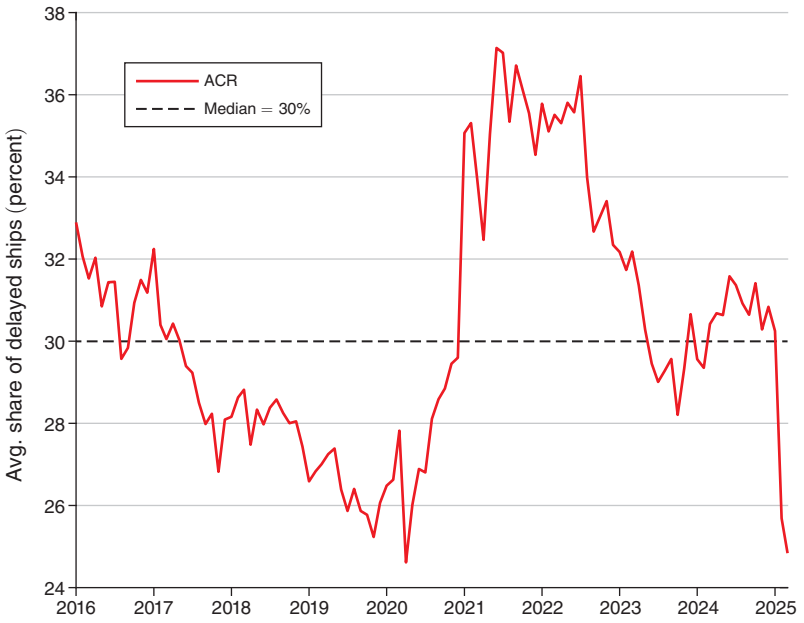


FIGURE 4. ACR INDEX

Notes: 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.

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 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:

$$ACR_t \equiv \sum_{p\in\mathcal{P}}\left[\frac{Delayed_{p,t} + Undelayed_{p,t}}{\sum_{p\in\mathcal{P}}\left(Delayed_{p,t} + Undelayed_{p,t}\right)} \cdot Congestion_{p,t}\right],$$

with $Congestion_{p,t}$ defined in equation (1).[14]

Figure 4 plots the ACR index. Before 2018, the index declined steadily and then stabilized around 28 percent, reaching a minimum of 25 percent 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

[14] We 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.

at 37 percent 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 percent) 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 percent. 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. 2026). By early 2025, however, the index dropped sharply below 25 percent. This decline reflects the gradual normalization of shipping routes, the easing of Panama Canal restrictions, and softer global trade volumes following newly imposed US tariffs that reduced import demand. In sum, the ACR index traces the sharp swings in global port congestion—from its COVID-19 peak, to postpandemic normalization, to renewed geopolitical pressures in 2024, and finally the marked easing in early 2025.

E. *Discussion*

Several aspects of our index merit further discussion, as they underpin the modeling assumptions in Section II, the identification restrictions in Section III, 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, Wang, Andersson, and Thun 2014).[15]

Supply chain disturbances (e.g., mobility restrictions, quarantines, and related disruptions discussed in Section IA) 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

[15] The 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 Supplemental Appendix A for details.

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 III, 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 Supplemental Appendix B and (ii) no statistical correlation between port congestion and oil price movements—reflecting capacity changes—in Supplemental Appendix A, despite the fact that higher oil prices directly incentivize slower steaming.

Alternative Indices.—Section ID compares our ACR index with other popular measures of global supply chain conditions, including the Harper Peterson 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, Smirnyagin, and Tsyvinski (2024).[17] Relative to these alternatives, supply chain shocks identified with the ACR produce impulse responses of the US 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.

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

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

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

Finally, integrating high-frequency AIS data with our spatial clustering algorithm allows us to construct congestion indices at frequencies higher than the monthly level. Supplemental 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.

II. 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 nontrivial. 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]

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.

A. *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.

[18] Supplemental Appendix D discusses the evidence of search frictions in the goods market and the relevance of transportation costs for the severance of commercial trade.

[19] We abstract from modeling producers' endogenous production decisions, as this would require a multicountry, multisector production network model to study the transmission of sectoral supply disruptions across countries (di Giovanni, Kalemli-Özcan, Silva, and Yildirim 2022, 2023).

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:

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

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

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

and the probability that a retailer meets a producer is

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

The function $f(\theta)$ satisfies $f(0) = 0$, $\lim_{\theta \to +\infty} f(\theta) = 1$, and $f_\theta(\theta) > 0$, while $q(\theta)$ satisfies $q(0) = 1$, $\lim_{\theta \to +\infty} q(\theta) = 0$, 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, $G(z) \equiv \Phi\left[\left(\ln z - \gamma\right)/\sigma\right]$, where $\Phi(\cdot)$ is the standard normal cdf.[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

[20] Our 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.

[21] The 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, Kalouptsidi, and Papageorgiou (2020); Bai and Li (2022); Dunn and Leibovici (2023); and Bai et al. (2026), but we keep the current setting for tractability.

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

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

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) shows that the present value of being a matched producer consists of the current profit margin, $\big(r(z) - z\big)l$, 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

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

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

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

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:

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

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

$$S(z) = X_M(z) - X_U + I_M(z) - I_U, \tag{7}$$

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

$$\eta\big(I_M(z) - I_U\big) = (1 - \eta)\big(X_M(z) - X_U\big). \tag{8}$$

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

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

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 ($\eta \to 0$), 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 cutoff transportation cost $\bar{z}$ above which matches become unprofitable and are therefore severed. This cutoff makes the total surplus in equation (7) equal to zero:

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

Substituting the value functions (3), (4), and (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:

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

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:

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

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:

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

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

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

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

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

B. *The Representative Household*

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

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

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:

$$pc + m \leq \mu + \underbrace{pc_s(\bar{z},\theta) - \rho l i_U - \textstyle\int_0^{\bar{z}} z' c_s(\bar{z},\theta)dG(z')}_{\text{Profits of producers and retailers}} + \underbrace{\textstyle\int_0^{\bar{z}} z' c_s(\bar{z},\theta)dG(z')}_{\text{Transportation costs}},$$

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:

$$\frac{\chi}{1+\chi}c^{-\frac{1}{\varepsilon}} = \frac{1}{1+\chi}\Big(\frac{m}{p}\Big)^{-\frac{1}{\varepsilon}}. \tag{15}$$

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

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

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.

C. *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 Supplemental Appendix E.5), we focus on the steady state. As shown in Section IID, 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 matched producers, given productive capacity l:

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

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,*

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

[22] We 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.

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 Supplemental 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—that is, 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*

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

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*:

(*i*) *The mapping*

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

is continuously differentiable and strictly increasing, *where*

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

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

(*ii*) $\lim_{p \to p_{\min}} c_s^{ss}(p) = 0$ *and* $\lim_{p \to +\infty} c_s^{ss}(p) = l$, *where* $p_{\min}$ *is defined as*

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

(*iii*) $c_s^{ss}(p)$ *is strictly increasing in* p *on* $[p_{\min}, +\infty)$ *and converges to a constant as* $p \to +\infty$;

(*iv*) $c_s^{ss}(p)$ *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 Supplemental Appendix E.2.

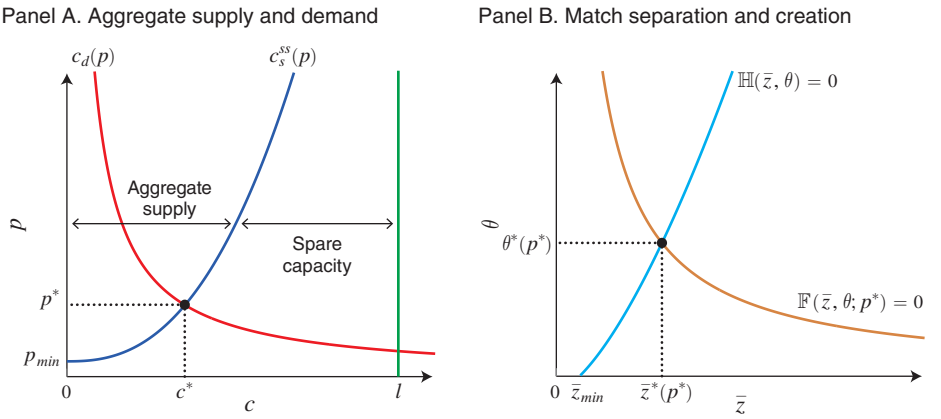


FIGURE 5. GRAPHICAL REPRESENTATION OF THE STEADY STATE

Notes: In panel A, $c_d(p)$, $c_s^{ss}(p)$, 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 B 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.

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 A of 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 B of 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 $(\bar{z}, \theta)$ plane.[23]

[23] For reference, Supplemental 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.

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

Adverse shock to:	Effects on: Consumption (or output) c	Price p	Reservation transportation cost $\bar{z}$	Product market tightness θ	Wholesale price r	Spare capacity $l - c$
Aggregate demand ($\mu \downarrow$ or $\chi \downarrow$)	$-$	$-$	$-$	$-$	$-$	$+$
Productive capacity ($l \downarrow$)	$-$	$+$	$+$	$+$	$+$	$-$
Assuming sticky $\bar{z}$ in short run, supply chain ($\gamma \uparrow$)	$-$	$+$	0	$-$	$\pm$	$+$

Notes: "0" denotes unchanged, "$+$" an increase, "$-$" a decrease, and "$\pm$" an undetermined effect. See Supplemental Appendix E.4 for the derivations of the comparative statics for each shock of interest.

D. *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 IE and justified in greater detail later in this section. Supplemental 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).

Aggregate Demand Shock.—Panel A of 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

[24] Comparable 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.

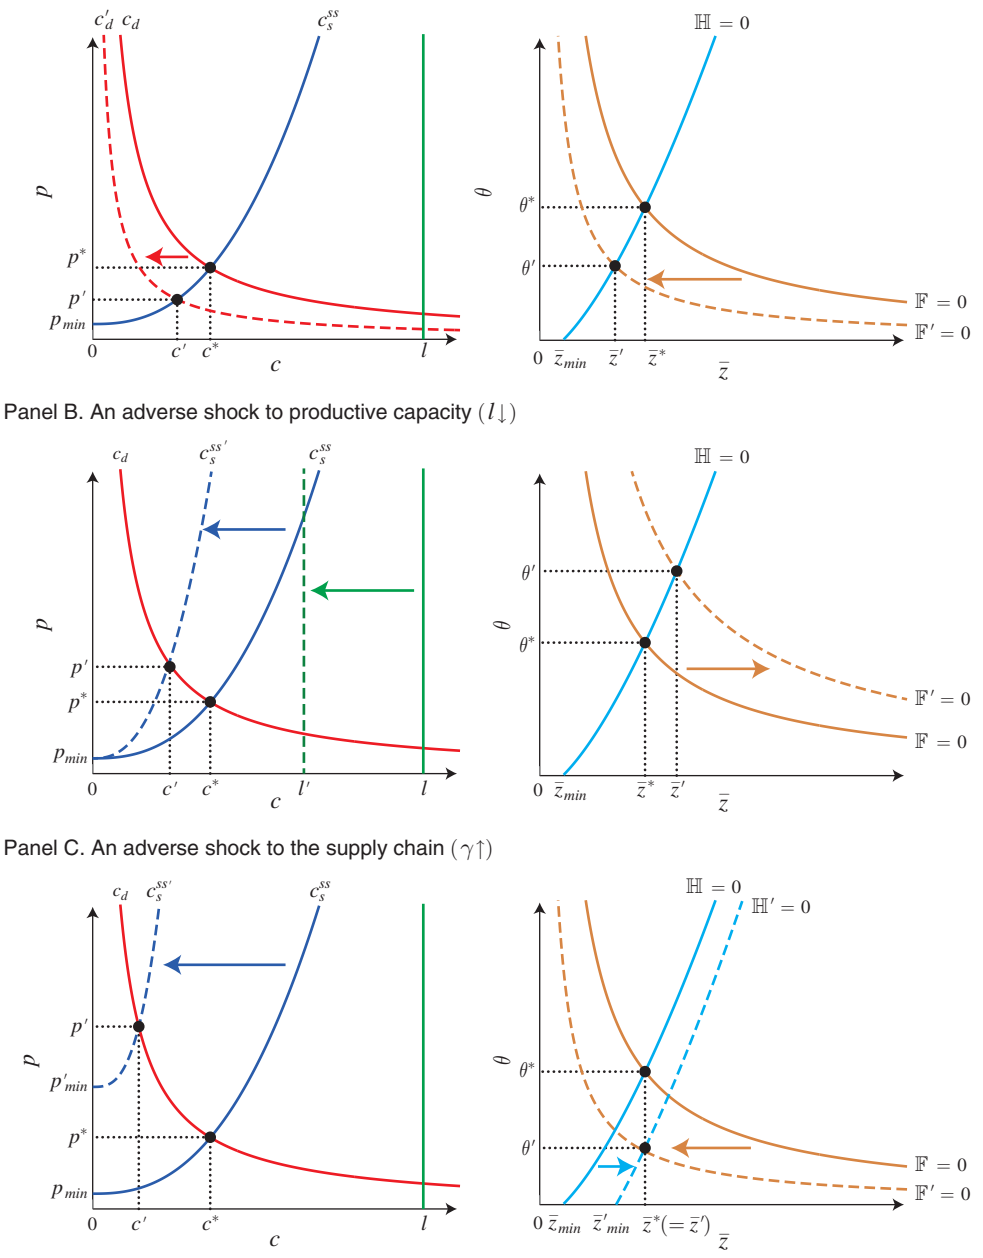


FIGURE 6. GRAPHICAL REPRESENTATION OF COMPARATIVE STATICS

of goods.[25] The resulting adjustment is a decline in the price required to clear the downstream retailer-household market.

[25] If 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.

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 spare productive capacity rises.

Productive Capacity Shock.—Panel B of 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 C of 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 IE, 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 C), 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 $\mathbb{F}_\gamma$ and $\mathbb{F}_\theta$ denote the partial derivatives of the match separation condition $\mathbb{F}(p, \bar{z}, \theta; \gamma) = 0$ with respect to γ and θ, respectively. Supplemental 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.

III. 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 to three distinct shocks, in line with the theoretical predictions in Table 1.[28]

[26] Supplemental Appendix 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.

[27] From 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.

[28] We 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.

A. *Empirical Framework*

Our empirical specification of the SVAR model follows Rubio-Ramírez, Waggoner, and Zha (2010) and Arias, Rubio-Ramírez, and Waggoner (2018):

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

where $\mathbf{y}_t$ is an $n \times 1$ vector of endogenous variables, $\boldsymbol{\omega}_t = \left[1, t\right]'$ 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 $\mathbf{y}_0, \ldots, \mathbf{y}_{1-L}$, 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

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

where $\mathbf{A}_+' = \left[\mathbf{A}_1' \ldots \mathbf{A}_L'\mathbf{C}'\right]$ and $\mathbf{x}_t' = \left[\mathbf{y}_{t-1}' \ldots \mathbf{y}_{t-L}'\boldsymbol{\omega}_t'\right]$. 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 $\mathbb{E}\left(\mathbf{u}_t\mathbf{u}_t'\right) = \boldsymbol{\Sigma} = \left(\mathbf{A}_0\mathbf{A}_0'\right)^{-1}$.

Motivated by the variables in our theoretical model, we estimate the baseline SVAR using monthly US data on real 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. Supplemental Appendix F provides a detailed overview of the external data and their sources.

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

$$SpareCapacityRate_t = \sum_{i \in \mathcal{C}}\left[\frac{Import_{i,2016}}{\sum_{i \in \mathcal{C}} Import_{i,2016}} \cdot \left(1 - CapacityUtilization_{i,t}\right)\right], \tag{23}$$

[29] "United States Trade," World Integrated Trade Solution, accessed October 15, 2024, https://wits.worldbank.org/CountrySnapshot/en/USA.

where $\mathcal{C} \equiv \{Mexico, Canada, China, Germany, Japan\}$, $Import_{i,2016}$ denotes US goods imports (customs basis) from country i in 2016, and $CapacityUtilization_{i,t}$ 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 US manufacturers' new orders and the number of unmatched producers (x_U) with the import-weighted average spare capacity of the specified US trading partners. Specifically,

$$Tightness_t = \frac{ManufactureNewOrder_t}{SpareCapacityDollar_t}. \tag{24}$$

The time series for US manufacturers' new orders ($ManufactureNewOrder_t$) is from Federal Reserve Economic Data (FRED). We construct the US dollar value of the import-weighted average spare capacity ($SpareCapacityDollar_t$) as follows:

$$SpareCapacityDollar_t = \sum_{i \in \mathcal{C}} \left[\frac{Import_{i,2016}}{\sum_{i \in \mathcal{C}} Import_{i,2016}} \cdot \left(\frac{IP_{i,t}}{CapacityUtilization_{i,t}} - IP_{i,t} \right) \right], \tag{25}$$

where $IP_{i,t}$ denotes total industrial production of country i in month t, measured in millions of constant 2005 US dollars. Finally, we use the import price as a proxy for the wholesale price to capture US 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.

We estimate the SVAR model using the Bayesian approach of Arias, Rubio-Ramírez, and Waggoner (2018); Arias, Caldara, and Rubio-Ramírez (2019); and Arias, Fernández-Villaverde, Rubio-Ramírez, and Shin (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 Supplemental 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

[30] We 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 $(\nu, \mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Omega})$. Parameters ν and $\mathbf{\Phi}$ govern the marginal prior distribution of $\text{vec}(\mathbf{A}_0)$, while $\mathbf{\Psi}$ and $\mathbf{\Omega}$ govern the prior distribution of $\text{vec}(\mathbf{A}_+)$ conditional on $\mathbf{A}_0$. We set $\nu = 0$, $\mathbf{\Phi} = \mathbf{0}_{n \times n}$, $\mathbf{\Psi} = \mathbf{0}_{m \times n}$, and $\mathbf{\Omega}^{-1} = \mathbf{0}_{m \times m}$. This parameterization yields prior densities equivalent to those in Uhlig (2005). Supplemental Appendix G.1 further shows that our results are robust to the prior-robust approach of Giacomini and Kitagawa (2021).

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 IE. 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.

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.

B. *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 percent and 90 percent 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]

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 percent, while spare capacity increases by slightly less than 0.5 percentage points. 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 percent, and the lower bound of the 90 percent probability band suggests that the effect could be as large as 3–4 percent; tightness then gradually converges back toward

[31] Two 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, Rubio-Ramírez, and Waggoner 2018).

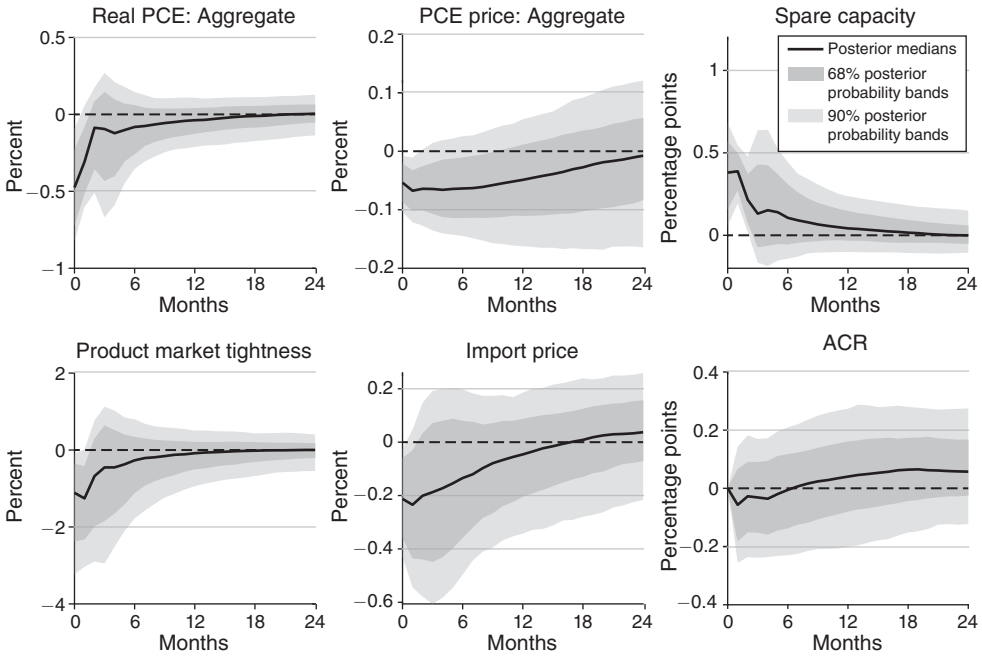


FIGURE 7. IRFS TO AN ADVERSE SHOCK TO AGGREGATE DEMAND

Note: The IRFs to a 1 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.

zero within two years. The PCE price index decreases modestly by about 0.1 percent 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 percent and remaining significantly below baseline through the first year. Last, 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 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 percent), 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 B 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 percent within one quarter and remain above baseline for six quarters, whereas PCE prices rise more modestly (peaking near 0.1 percent) and fade more slowly. Real PCE reverses quickly from its initial dip, rising by about 0.5–0.6 percent within the first quarter before mean reverting; by contrast, spare capacity troughs just above -0.5 percentage points around the one-quarter mark and then recovers gradually toward zero. The ACR index is muted

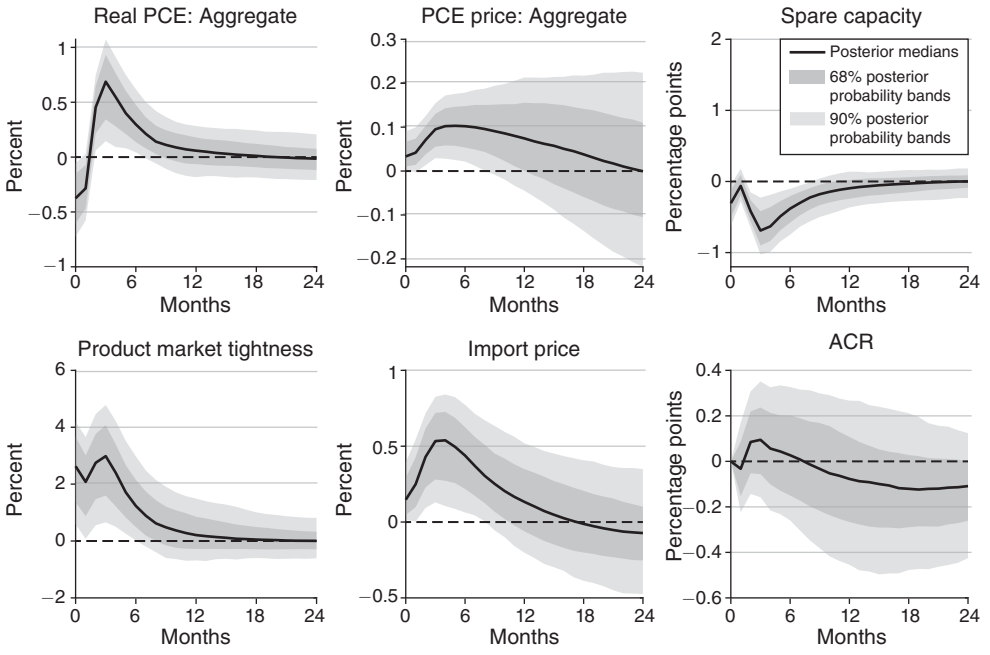


FIGURE 8. IRFS TO AN ADVERSE SHOCK TO PRODUCTIVE CAPACITY

Note: The IRFs to a 1 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.

on impact, climbs to roughly 0.1 percentage points at about three months, and then drifts below zero around the three-quarter horizon, stabilizing near −0.1 percentage points, suggesting complementarities between production constraints and transportation frictions in the transmission of capacity shocks.

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 percent around one year, and the lower bound of the 90 percent probability band stays above zero throughout, indicating a robust inflationary effect.

Import prices, which are unrestricted in our specification, also rise persistently. The 68 percent probability band remains above zero across most horizons, and the median peaks above 0.3 percent 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 percent), turns positive within four to six months (peaking near 0.2 percent), and then gradually mean reverts. Spare capacity rises initially (0.3 percentage points) 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 percentage points) and then decays gradually,

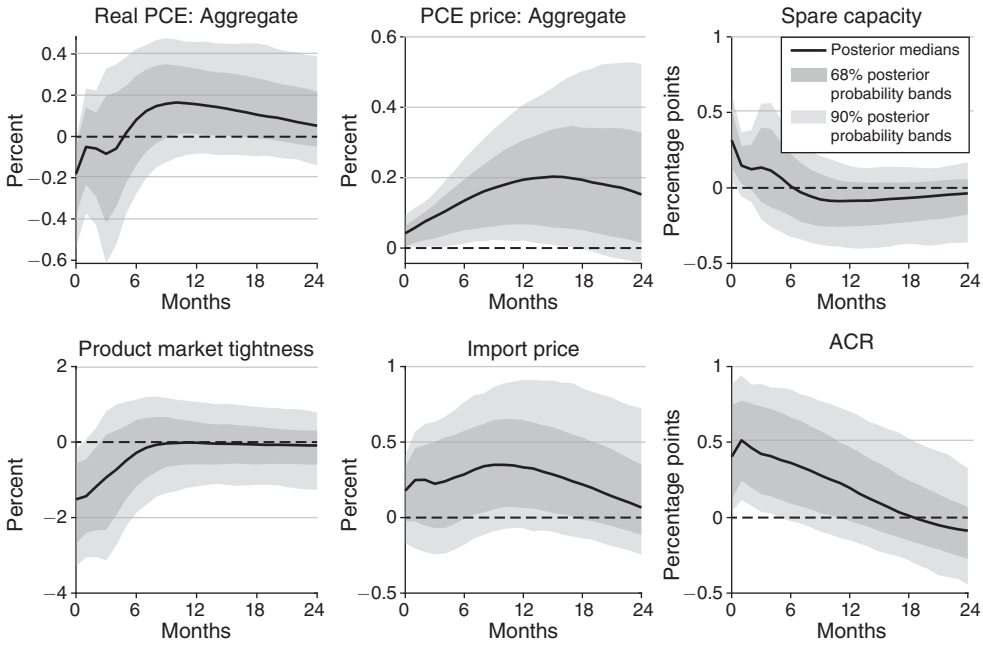


FIGURE 9. IRFS TO AN ADVERSE SHOCK TO THE SUPPLY CHAIN

Note: The IRFs to a 1 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.

crossing zero around the 18-month horizon, consistent with congestion easing as logistics adjust.

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 percent at short horizons to about 20–25 percent by 6 to 12 months and remains the largest contributor thereafter, while the contributions of demand and capacity shocks decline toward 5–10 percent. For import prices, productive capacity shocks lead early (peaking near 25 percent around six months), but the supply chain share increases steadily and eventually overtakes the capacity share, reaching roughly 15 percent 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 percent at three to six months and remain the largest contributor even at two years, while the demand share is front-loaded—high on impact but falling toward 10 percent—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 percent 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 percent across horizons—while demand and capacity shocks

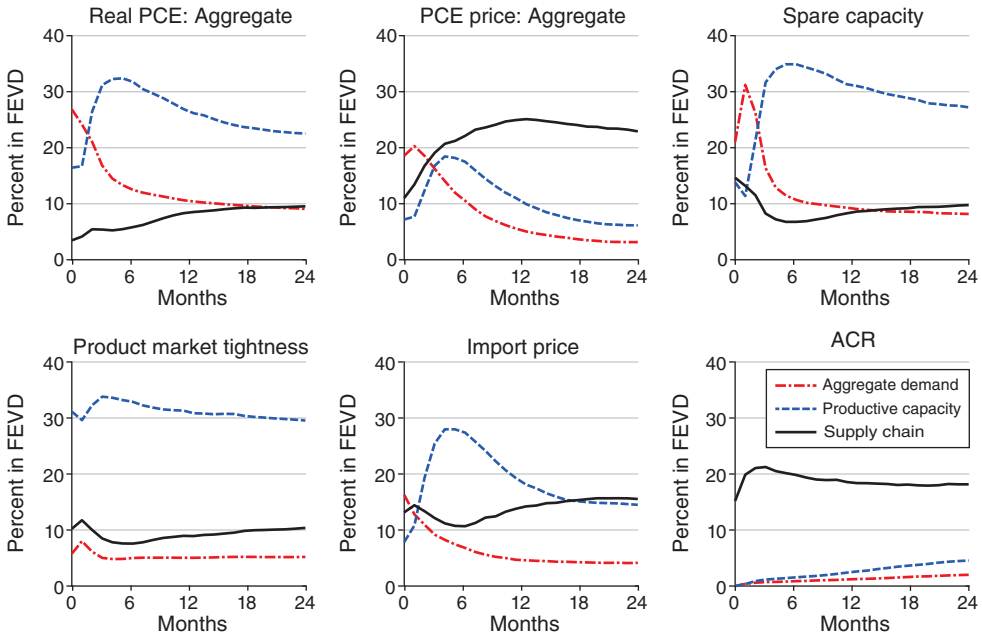


FIGURE 10. FORECAST ERROR VARIANCE DECOMPOSITIONS (FEVDS) FROM THE SVAR

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).

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 IE.[32]

Figure 11 presents the key empirical result of our analysis: the cumulative historical contributions of the three identified shocks to US half-on-half (HOH) headline PCE inflation over the period September 2016 to March 2025.[33] The main findings from the decomposition can be summarized in five points:

(i) Before the pandemic (2016–2019), contributions from all three shocks were modest and often offsetting, with the supply chain component generally slightly negative.[34]

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

[33] Historical 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 (US HOH headline PCE inflation).

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

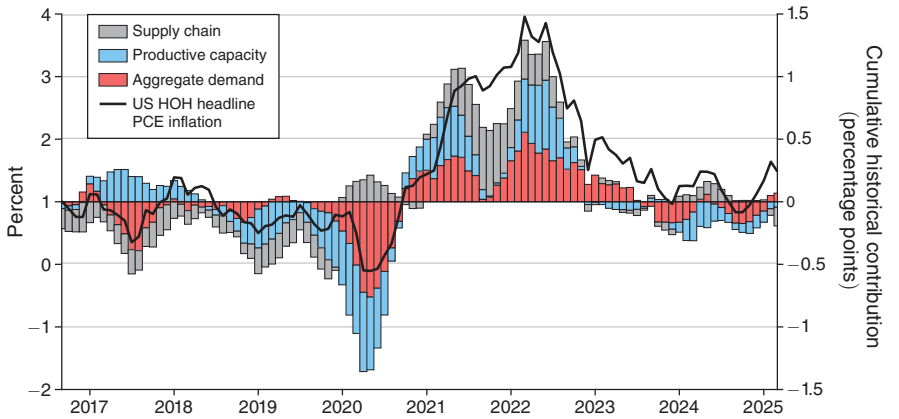


FIGURE 11. HISTORICAL DECOMPOSITION OF US HALF-ON-HALF (HOH) HEADLINE PCE INFLATION

Notes: The solid line plots US headline PCE inflation, measured as HOH growth of the PCE chain-type price index, against the left axis (percent). The shaded bars plot the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (percentage points). 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.

(ii) 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.

(iii) 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 US trading partners that constrained effective capacity.

(iv) 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.

(v) 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. (2026) for details), while demand and capacity continued to exert a modest drag.[35]

[35] Two 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; (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.

C. *Robustness Checks and Extensions*

We conduct a range of robustness checks on our baseline results beyond those previously discussed. First, Supplemental Appendix G.2 reestimates 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 forecast error variance decomposition (FEVD) rankings remain stable. In addition, the shape 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, Supplemental 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 attenuates the supply chain price effects while leaving demand- and capacity-driven IRFs essentially unchanged. Because the baseline six-variable model may absorb omitted sectoral substitution and oil shocks, it provides an upper-bound estimate of supply chain price pass-through, whereas this augmented model serves as a conservative lower bound.

Third, Supplemental 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., Baqaee 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, Supplemental Appendix G.5 shows that reestimating 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, Supplemental 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, Fernández-Villaverde, Li, and Zanetti 2026b), we analyze the effectiveness of monetary policy during global supply chain disruptions. Guided by the theoretical predictions of our model in Section II 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 reestimate the US 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, Lorenzoni, Straub, and Werning 2022; Ferrante, Graves, and Iacoviello 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

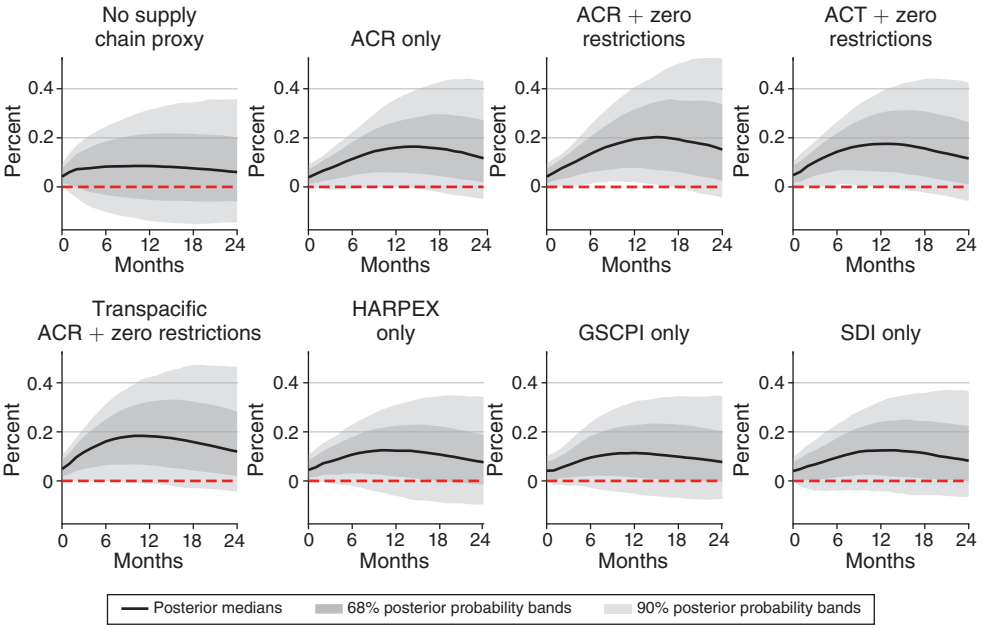


FIGURE 12. PRICE RESPONSES TO AN ADVERSE SUPPLY CHAIN SHOCK ACROSS ALTERNATIVE PROXIES

Notes: IRFs of the US PCE chain-type price index to a 1 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) transpacific ACR with the same zero restrictions; (vi) HARPEX; (vii) the New York Fed's GSCPI (Benigno, di Giovanni, Groen, and Noble 2022); and (viii) the SDI from Smirnyagin and Tsyvinski (2022) and Liu, Smirnyagin, and Tsyvinski (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.

and sector-specific dynamics. For the euro area, an analogous SVAR points to a dominant and persistent role for oil price shocks in 2021–2022, reflecting Europe's heightened energy dependence. Full results are available upon request.

D. *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 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 transpacific ACR—i.e., a targeted ACR index for major ports along the transpacific 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 IE—further sharpens identification. The resulting price responses exhibit higher peaks, and the lower bound of the 90 percent probability band diverges from

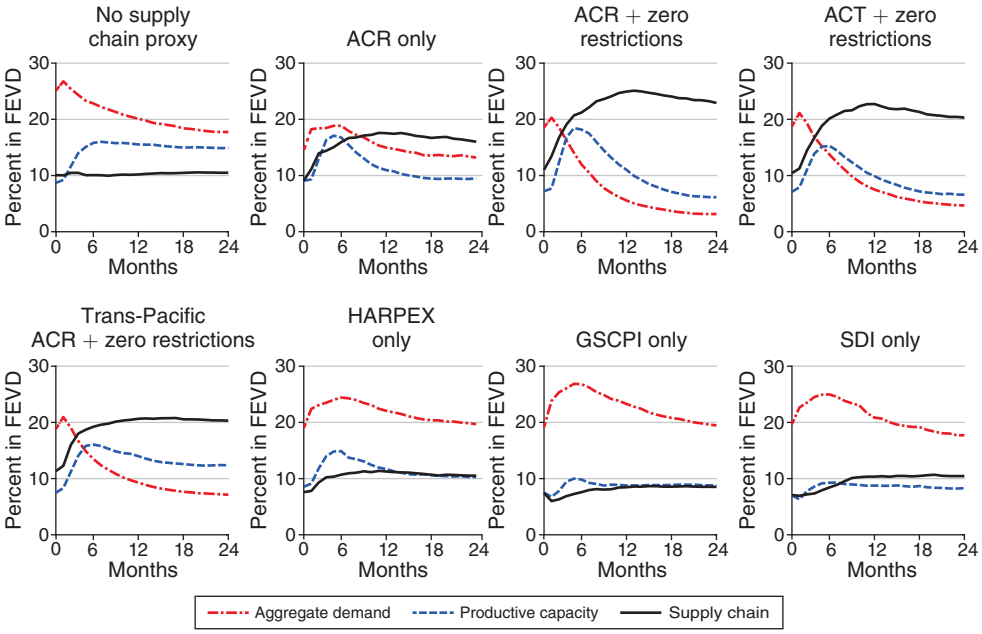


FIGURE 13. FEVD OF THE US PCE PRICE INDEX ACROSS ALTERNATIVE PROXIES

Note: The figure plots the posterior median shares of the forecast error variance of the US PCE chain-type price index attributable to aggregate demand, productive capacity, and supply chain shocks across horizons, for each of the eight specifications illustrated in Figure 12.

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.

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 Supplemental Appendix H.

IV. 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, Mandelman, Yu, and Zanetti 2021, 2025; Fernández-Villaverde, Yu, and Zanetti 2024). 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, Yu, and Zanetti 2025). We plan to explore some of these extensions in future work.

REFERENCES

Acharya, Viral V., Matteo Crosignani, Tim Eisert, and Christian Eufinger. 2023. "How Do Supply Shocks to Inflation Generalize? Evidence from the Pandemic Era in Europe." NBER Working Paper 31790.

Alessandria, George, Shafaat Yar Khan, Armen Khederlarian, Carter Mix, and Kim J. Ruhl. 2023. "The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020–2022." *Journal of International Economics* 146: 103788.

Allen, Treb, and Costas Arkolakis. 2014. "Trade and the Topography of the Spatial Economy." *Quarterly Journal of Economics* 129 (3): 1085–140.

Antràs, Pol. 2023. "An 'Austrian' Model of Global Value Chains." NBER Working Paper 30901.

Arias, Jonas E., Dario Caldara, and Juan F. Rubio-Ramírez. 2019. "The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure." *Journal of Monetary Economics* 101: 1–13.

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

Arias, Jonas E., Juan F. Rubio-Ramírez, and Daniel F. Waggoner. 2018. "Inference Based on Structural Vector Autoregressions Identified with Sign and Zero Restrictions: Theory and Applications." *Econometrica* 86 (2): 685–720.

Ascari, Guido, Dennis Bonam, and Andra Smadu. 2024. "Global Supply Chain Pressures, Inflation, and Implications for Monetary Policy." *Journal of International Money and Finance* 142: 103029.

Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, Ricardo Marto, and Francesco Zanetti. 2026. "Shipping to America." Unpublished.

Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti. 2026a. *Data and Code for: "The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory."* American Economic Association; distributed by Inter-university Consortium for Political and Social Research. https://doi.org/10.3886/E247539V1.

Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti. 2026b. "State Dependence of Monetary Policy during Global Supply Chain Disruptions." Unpublished.

Bai, Xiwen, and Yiliang Li. 2022. "The Congestion Effect of Oil Transportation and Its Trade Implications." Unpublished.

Bai, Xiwen, Zhongjun Ma, Yao Hou, Yiliang Li, and Dong Yang. 2023. "A Data-Driven Iterative Multi-attribute Clustering Algorithm and Its Application in Port Congestion Estimation." *IEEE Transactions on Intelligent Transportation Systems* 24 (11): 12026–37.

Baqaee, David, and Emmanuel Farhi. 2019. "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem." *Econometrica* 87 (4): 1155–203.

Baqaee, David, and Emmanuel Farhi. 2020. "Nonlinear Production Networks with an Application to the COVID-19 Crisis." NBER Working Paper 27281.

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

Bauer, Michael D., and Eric T. Swanson. 2023. "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual* 37: 87–155.

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

Benigno, Pierpaolo, and Gauti B. Eggertsson. 2023. "It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-linear Phillips Curve." NBER Working Paper 31197.

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

Binetti, Alberto, Francesco Nuzzi, and Stefanie Stantcheva. 2024. "People's Understanding of Inflation." *Journal of Monetary Economics* 148: 103652.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. "Geography, Transportation, and Endogenous Trade Costs." *Econometrica* 88 (2): 657–91.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2024. "Investment in Infrastructure and Trade: The Case of Ports." NBER Working Paper 32503.

Brancaccio, Giulia, Myrto Kalouptsidi, Theodore Papageorgiou, and Nicola Rosaia. 2023. "Search Frictions and Efficiency in Decentralized Transport Markets." *Quarterly Journal of Economics* 138 (4): 2451–503.

Bureau of Transportation Statistics. 2021. *On National Maritime Day and Every Day, U.S. Economy Relies on Waterborne Shipping*. Data Spotlights. https://www.bts.gov/data-spotlight/national-maritime-day-and-every-day-us-economy-relies-waterborne-shipping (accessed June 17, 2024).

Cerdeiro, Diego A., and Andras Komaromi. 2020. "Supply Spillovers during the Pandemic: Evidence from High-Frequency Shipping Data." IMF Working Paper 2020/284.

Chen, Longbiao, Daqing Zhang, Xiaojuan Ma, Leye Wang, Shijian Li, Zhaohui Wu, and Gang Pan. 2016. "Container Port Performance Measurement and Comparison Leveraging Ship GPS Traces and Maritime Open Data." *IEEE Transactions on Intelligent Transportation Systems* 17 (5): 1227–42.

Comín, Diego A., Robert C. Johnson, and Callum J. Jones. 2023. "Supply Chain Constraints and Inflation." Federal Reserve Board Finance and Economics Discussion Series 2023-075.

di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A. Yildirim. 2022. "Global Supply Chain Pressures, International Trade, and Inflation." NBER Working Paper 30240.

di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A. Yildirim. 2023. "Pandemic-Era Inflation Drivers and Global Spillovers." NBER Working Paper 31887.

Dunn, Jason, and Fernando Leibovici. 2023. "Navigating the Waves of Global Shipping: Drivers and Aggregate Implications." Unpublished.

Fan, Lei, William W. Wilson, and Bruce Dahl. 2012. "Congestion, Port Expansion and Spatial Competition for US Container Imports." *Transportation Research Part E: Logistics and Transportation Review* 48 (6): 1121–36.

Feng, Mingxiang, Shih-Lung Shaw, Guojun Peng, and Zhixiang Fang. 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, Jesús, Federico Mandelman, Yang Yu, and Francesco Zanetti. 2021. "The 'Matthew Effect' and Market Concentration: Search Complementarities and Monopsony Power." *Journal of Monetary Economics* 121: 62–90.

Fernández-Villaverde, Jesús, Federico Mandelman, Yang Yu, and Francesco Zanetti. 2025. "Search Complementarities, Aggregate Fluctuations, and Fiscal Policy." *Review of Economic Studies* 92 (4): 2502–36.

Fernández-Villaverde, Jesús, Yang Yu, and Francesco Zanetti. 2024. "Technological Synergies, Heterogeneous Firms, and Idiosyncratic Volatility." NBER Working Paper 32247.

Ferrante, Francesco, Sebastian Graves, and Matteo Iacoviello. 2023. "The Inflationary Effects of Sectoral Reallocation." *Journal of Monetary Economics* 140: S64–81.

Finck, David, and Peter Tillmann. 2022. "The Macroeconomic Effects of Global Supply Chain Disruptions." Bank of Finland Institute for Emerging Economies Discussion Paper 14/2022.

Finck, David, Mathias Klein, and Peter Tillmann. 2024. "The Inflationary Effects of Global Supply Chain Shocks: Evidence from Swedish Microdata." Unpublished.

Freightos. 2024. *Shipping from Shanghai, Shanghai to Los Angeles, CA: Air, Sea & Container Freight*. Freightos.

Fuchs, Simon, and Woan Foong Wong. 2022. "Multimodal Transport Networks." Unpublished.

Ganapati, Sharat, Woan Foong Wong, and Oren Ziv. 2024. "Entrepôt: Hubs, Scale, and Trade Costs." *American Economic Journal: Macroeconomics* 16 (4): 239–78.

Ghassibe, Mishel. 2024. "Endogenous Production Networks and Non-linear Monetary Transmission." Unpublished.

Ghassibe, Mishel, and Francesco Zanetti. 2022. "State Dependence of Fiscal Multipliers: The Source of Fluctuations Matters." *Journal of Monetary Economics* 132: 1–23.

Giacomini, Raffaella, and Toru Kitagawa. 2021. "Robust Bayesian Inference for Set-Identified Models." *Econometrica* 89 (4): 1519–56.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2022. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" *American Economic Review* 112 (5): 1437–74.

Harding, Martín, Jesper Lindé, and Mathias Trabandt. 2023. "Understanding Post-COVID Inflation Dynamics." *Journal of Monetary Economics* 140: S101–18.

Heiland, Inga, Andreas Moxnes, Karen Helene Ulltveit-Moe, and Yuan Zi. Forthcoming. "Trade from Space: Shipping Networks and the Global Implications of Local Shocks." *Review of Economics and Statistics*.

Karimi-Mamaghan, Maryam, Mehrdad Mohammadi, Amir Pirayesh, Amir Mohammad Karimi-Mamaghan, and Hassan Irani. 2020. "Hub-and-Spoke Network Design under Congestion: A Learning Based Metaheuristic." *Transportation Research Part E: Logistics and Transportation Review* 142: 102069.

Kasahara, Hiroyuki, and Beverly Lapham. 2013. "Productivity and the Decision to Import and Export: Theory and Evidence." *Journal of International Economics* 89 (2): 297–316.

Khalil, Makram, and Marc-Daniel Weber. 2022. "Chinese Supply Chain Shocks." Deutsche Bundesbank Discussion Paper 44/2022.

Kim, Se-Jik, and Hyun Song Shin. 2023. "Theory of Supply Chains: A Working Capital Approach." Bank for International Settlements Working Paper 1070.

Komaromi, Andras, Diego A. Cerdeiro, and Yang Liu. 2022. "Supply Chains and Port Congestion around the World." IMF Working Paper 2022/59.

Liu, Ernest, Vladimir Smirnyagin, and Aleh Tsyvinski. 2024. "Supply Chain Disruptions and Supplier Capital in U.S. Firms." Unpublished.

Li, Yiliang, Xiwen Bai, Qi Wang, and Zhongjun Ma. 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.

Meng, Qiang, Shuaian Wang, Henrik Andersson, and Kristian Thun. 2014. "Containership Routing and Scheduling in Liner Shipping: Overview and Future Research Directions." *Transportation Science* 48 (2): 265–80.

Merendino, Alfonso, and Tommaso Monacelli. 2025. "Supply Chain Uncertainty, Energy Prices, and Inflation." Unpublished.

Michaillat, Pascal, and Emmanuel Saez. 2015. "Aggregate Demand, Idle Time, and Unemployment." *Quarterly Journal of Economics* 130 (2): 507–69.

Michaillat, Pascal, and Emmanuel Saez. 2022. "An Economical Business-Cycle Model." *Oxford Economic Papers* 74 (2): 382–411.

Mumtaz, Haroon, and Francesco Zanetti. 2012. "Neutral Technology Shocks and the Dynamics of Labor Input: Results from an Agnostic Identification." *International Economic Review* 53 (1): 235–54.

Mumtaz, Haroon, and Francesco Zanetti. 2015. "Labor Market Dynamics: A Time-Varying Analysis." *Oxford Bulletin of Economics and Statistics* 77 (3): 319–38.

Notteboom, Theo, Athanasios Pallis, and Jean-Paul Rodrigue. 2022. *Port Economics, Management and Policy*. Routledge.

Planet Team. 2020. *Planet Application Program Interface: In Space for Life on Earth.* San Francisco, CA: Planet Labs. https://api.planet.com.

Rubio-Ramírez, Juan F., Daniel F. Waggoner, and Tao Zha. 2010. "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference." *Review of Economic Studies* 77 (2): 665–96.

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

Smirnyagin, Vladimir, and Aleh Tsyvinski. 2022. "Macroeconomic and Asset Pricing Effects of Supply Chain Disasters." NBER Working Paper 30503.

Stopford, Martin. 2008. *Maritime Economics*. 3rd ed. Routledge.

Talley, Wayne K. 2009. *Port Economics*. Routledge.

Talley, Wayne K., and Man Wo Ng. 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*. National Academies Press.

Uhlig, Harald. 2005. "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics* 52 (2): 381–419.

UNCTAD (**United Nations Conference on Trade and Development**)**.** 2019. *Review of Maritime Transport 2019.* United Nations.

Xu, Le, Yang Yu, and Francesco Zanetti. 2025. "The Adoption and Termination of Suppliers over the Business Cycle." *Journal of Monetary Economics* 151: 103730.