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

Xiwen Bai† Jesús Fernández-Villaverde‡ Yiliang Li§

Francesco Zanetti¶

June 19, 2024

Abstract

We study the causal effects and policy implications of global supply chain disruptions. We construct a new index that measures the state of the global supply chain from the mandatory automatic identification system data of container ships and propose a novel spatial clustering algorithm that determines real-time congestion from the positions, speeds, and headings of container ships in major ports around the globe. We develop a model with search frictions between producers and retailers that links spare productive capacity with congestion in the goods market and the responses of output and prices to supply chain shocks. The co-movements of output, prices, and spare capacity yield unique identification restrictions that allow us to study the causal effects on macroeconomic outcomes. We document how supply chain shocks drove U.S. inflation during 2021 but that, from 2022 onward, traditional demand and supply shocks also played an important role in explaining inflation. Finally, we show how monetary policy is more effective in taming inflation amid supply chain disruptions than in regular circumstances.

JEL Classification: E32, E58, J64.
Keywords: Supply chain disruptions, search-and-matching in the goods market, SVAR, state-dependence of monetary policy.

*We are grateful to Klaus Adam, Hassan Afrouzi, Fernando Álvarez, Gianluca Benigno, Hilde Bjørnland, Dennis Bonam, Bjoern Bruegemann, Vasco Carvalho, Diego Comín, Xiaomin Cui, Elena María Díaz, Julian di Giovanni, Gauti Eggertsson, Mariassunta Giannetti, Mishel Ghassibe, Sebastian Heise, Bo Hu, Callum Jones, Şebnem Kalemli-Özcan, Enisse Kharroubi, Nobuhiro Kiyotaki, Laura Lebastard, Ernest Liu, Marco Lombardi, Emi Nakamura, George Nikolakoudis, Ekaterina Peneva, Theodore Papageorgiou, Giorgio Primiceri, Omar Rachedi, Stephen Redding, Ricardo Reis, Esteban Rossi-Hansberg, Katheryn Russ, Karthik Sastry, Adam Shapiro, Hyun Song Shin, Frank Smets, Vladimir Smirnyagin, Bo Sun, Aleh Tsyvinski, Harald Uhlig, Cliff Winston, Christian Wolf, Le Xu, Li Yu, Yang Yu, Penghui Yin, Bianca Zanetti, Charles Zhang, Yuan Zi, and participants at numerous conferences and seminars for their comments and suggestions. Zhongjun Ma provided excellent research assistance. Francesco Zanetti gratefully acknowledges financial support from the British Academy.
† Bai: Tsinghua University, China. xiwenbai@mail.tsinghua.edu.cn. ‡ Fernández-Villaverde: University of Pennsylvania, U.S. jesusfv@econ.upenn.edu. § Li: University of International Business and Economics, China. yiliang_li@uibe.edu.cn. ¶ Zanetti: University of Oxford, U.K. francesco.zanetti@economics.ox.ac.uk.

1. Introduction

The world economy is organized around an intricate global supply chain. Any sudden and large shock to this global supply chain, such as those triggered by a war, the COVID-19 pandemic, or the ongoing Red Sea crisis, might have large consequences for output, inflation, and spare productive capacity. Furthermore, global supply chain shocks might also shift the trade-offs that policymakers face when stabilizing the economy.[1]

However, measuring the causal effects of a global supply chain shock and, hence, being able to design optimal policy responses to it, is challenging. First, researchers need to measure the state of the global supply chain. Existing indices are often inferred from changes in shipping prices or information from surveys on potential disruptions gleaned from the Purchasing Managers' Index (PMI). These measures are problematic. Shipping prices reflect movements in the demand for tradable goods, market expectations, or fuel prices that are unrelated to supply chain disruptions. Surveys of managers are subject to potentially large measurement errors arising from the subjective perceptions of interviewees. Instead, the ideal measurement requires accurate data that tracks the flow of goods around the globe.

Second, researchers need a theoretical framework to derive the identification assumptions necessary for the causality analysis of supply chain disruption shocks. At any given moment, the global supply chain is driven by a combination of aggregate demand, aggregate supply, and supply chain shocks. These shocks can only be disentangled by applying identification restrictions derived from theory. However, there is currently no standard model that simultaneously addresses the increase in spare productive capacity alongside the shortage of goods and the scarcity of supply in the retail market, which, as we argue later, are crucial for differentiating supply chain disturbances from other shocks.

Our paper addresses these issues by developing (i) a new index of global supply chain disruptions that tracks the congestion of container ships at major ports worldwide using high-frequency maritime satellite data and that provides a real-time and accurate measure of the state of the global supply chain, and (ii) a novel theory that accounts for the coexistence of elevated spare capacity for producers and scarcity of supply in the retail market, and examines how these factors

[1] Many articles in the popular press and policy institutions discuss the relevance of disruptions to the supply chain for economic performance and policy. See, for instance, Attinasi et al. (2021), The White House (2021), Lane (2022), and The World Bank (2022).

affect the responses of output and prices during a disruption to the supply chain. Leveraging the identification assumptions derived from our theoretical framework and a Bayesian structural vector autoregression (SVAR), we can separate the various shocks driving our index of global supply chain disruptions and measure the causal effects of such disruptions on aggregate economic outcomes. For example, if we observe an increase in our index, we will be able to tell how much of this increase was driven by each of the present and past values of aggregate demand, aggregate supply, and supply chain shocks and how each shock contributed to the dynamics of output and inflation.

Relevance for the future. The importance of the answers to points (i) and (ii) above is that it is likely that the world economy might experience again large disruptions to the global supply chain, such as those triggered by wars, geostrategic realignments, blockades, sanctions, or even another pandemic. Far from being just a postmortem of what happened during the COVID-19 pandemic, our analysis distills important lessons for the future.

Measuring the state of the global supply chain. We follow maritime economics by assessing the health of the global supply chain through the degree of congestion at container ports worldwide. This concept has gained widespread recognition. As early as 2006, Transportation Research Board Executive Committee pinpointed congestion as a critical issue impacting all modes of transportation and logistical functions. This perspective was further corroborated by the influential work of Fan et al. (2012), and more recently, Brancaccio et al. (2024), who documented the impact of port congestion on the efficiency and reliability of global supply chains.

The reason is that container shipments play a pivotal role in global trade. Around 60% of the total value of world seaborne trade passes through container ports (Coşar and Demir, 2018; UNCTAD, 2019; OECD and EUIPO, 2021), implying that even a mild increase in port congestion can generate large imbalances between the supply and demand for tradable goods, thereby putting enormous strains on the global supply chain.

Our port congestion index offers a precise, real-time assessment of the regular flow of goods worldwide. The high-frequency maritime satellite data enable us to track the movements of container ships with virtually no measurement error. Furthermore, port congestion provides a critical advantage in identifying the causal effects of supply chain disruption shocks. While port congestion is largely determined by the number of ship visits, container ships operate on fixed schedules that only respond to fluctuations in demand or productive capacity with a significant

lag. Short-term amendments to such schedules or other contractual terms are rare because they involve severe penalties and switching costs.[2] Consequently, as we will describe below, our identification strategy only requires that schedules or contracts are fixed for at least one month, ensuring that our port congestion index does not respond to demand or capacity shocks within a month's period. Even if only because travel time for a ship between ports often takes more than one month, we are confident our identification assumption is valid.

We quantify port congestion around the globe from 2017 to 2023 using granular shipping data from the automatic identification system (AIS), the long-range identification and tracking system on container ships mandated by the International Maritime Organization (IMO), the specialized agency of the United Nations responsible for regulating the shipping industry worldwide. By developing a novel machine learning-based spatial clustering algorithm that utilizes the positions, speeds, and headings of container ships recorded in the AIS data, we construct a new dataset that provides a measure of congestion at individual ports, which we then aggregate across ports to develop the first high-frequency index of the average congestion rate (ACR).

Our index indicates that strains on the global supply chain and the escalation of port congestion related to the COVID-19 pandemic began in the second half of 2020 and remained elevated until the second half of 2022. It shows that the average proportion of container ships experiencing delays in their loading and unloading operations upon arrival at ports increased from 25% to 37%. At the same time, the average duration of such delays rose from 5.5 to 13.5 hours. Combined with the fact that nearly 80% of world trade is shipped indirectly, and the average shipment stops at five additional ports before reaching its destination (Ganapati et al., 2021), our ACR index documents the large obstructions to the systematic flows of container ships around the globe during the COVID-19 pandemic.

Theoretical framework. Next, we develop a model that accounts for the imbalances between supply and demand for goods resulting from supply chain disturbances. Our model is built around the search and matching frictions between producers and retailers, each based in different locations. Also, the shipments of goods to retailers require producers to pay transportation costs.

Our model is inspired by the literature on disequilibrium models from the 1970s (e.g., Barro and Grossman 1971), but recast in a microfounded framework with search and matching frictions

[2]The terms of shipping services often extend beyond one year, and the design of shipping routes, service schedules, and speeds are revised every three to six months on average (Stopford, 2008; Meng et al., 2014).

by Michaillat and Saez (2015, 2022) and Ghassibe and Zanetti (2022). By separating producers and retailers and incorporating transportation costs, we can jointly generate spare capacity for producers, scarcity of supply, and increased congestion in the retail market.

The presence of search frictions introduces trading externalities that limit the allocative role of prices: retailers and producers of goods face a probability of failing to match with each other. In other words, our framework accounts for rationing in the retail market that price adjustments cannot eliminate. Instead, trading is determined by the relative number of retailers and producers, which is influenced by supply chain disruptions.

We assume that a supply chain shock can take two alternative forms. The first is an increase in transportation costs, which is motivated by the evidence showing that ships were tied up at ports due to increased congestion, leading to a significant shortage in shipping supply and soaring shipping prices during the COVID-19 pandemic (Alessandria et al., 2023; Dunn and Leibovici, 2023). The rise in transportation costs reduces the number of profitable shipments and curtails the volume of shipped goods, leading to a fall in the goods supply available to retailers as well as increasing the spare capacity for producers and prices.

Second, the supply chain disruption can also be modeled as a reduction in matching efficiency between producers and retailers. This modeling choice is consistent with the difficulties in establishing new partnerships when established trading relationships are disrupted by acute input delays, such as those experienced during the pandemic (Smirnyagin and Tsyvinski, 2022; Liu et al., 2024). A reduction in matching efficiency decreases the probability of retailers meeting producers, thus imposing larger costs on retailers to form a match. We show that both alternative modelings of the supply chain disruption deliver the same predictions regarding the effects on consumption, prices, and spare capacity.

Our model demonstrates that the responses of macro aggregates to a supply chain disruption shock differ from those to standard demand and supply shocks. Unlike demand shocks, disruptions to the supply chain result in negative co-movements between output and prices. Although traditional supply shocks – in the form of adverse shocks to productive capacity – are also characterized by negative co-movements between output and prices, disruptions to the supply chain uniquely increase spare capacity for producers due to the reduction in the shipments of goods, whereas traditional supply shocks decrease it. Such a difference is intuitive: supply chain disruptions do not alter productive capacity in the goods market but impede the flow of goods to

retailers, leading to increased spare capacity and a deficient supply in the retail market. Thus, the rise in spare capacity, along with higher prices and lower output, enables the identification of supply chain disturbances.

The causal effects of supply chain disruptions. We apply our theoretical predictions on the responses of the endogenous variables to identify a Bayesian SVAR enriched with our ACR index. Our main empirical results are as follows.

First, a disruption shock to the supply chain leads to a large and immediate drop in U.S. real GDP and a surge in spare capacity, proxied by an import-weighted average spare capacity rate of the top five exporting countries to the U.S. – namely, Mexico, Canada, China, Germany, and Japan, which collectively account for more than 53% of goods imported into the U.S. Additionally, the supply chain disruption shock generates a persistent positive response of U.S. goods inflation, an observation consistent with recent empirical evidence (Bekaert et al., 2020; Benigno et al., 2022; Finck and Tillmann, 2022; De Santis, 2024). As predicted by our model, the productive capacity shock and the supply chain shock differ in their effects on spare capacity. For the capacity shock, spare capacity consistently falls, while for the supply chain shock, it persistently increases, with the median response reverting to zero at the one-quarter mark.

Second, the historical decomposition shows that U.S. goods inflation since 2020 has experienced four phases. In the first phase (2020), the sharp fall in inflation was mainly driven by a significant contraction in aggregate demand coinciding with the first wave of the COVID-19 pandemic across the world. In the second phase (2021), inflation was largely caused by global supply chain disruptions. In the third phase (2022), adverse shocks to productive capacity kept inflation elevated. In the final phase (2023), inflation subsided due to a combination of weakened demand, strengthened capacity, and supply chain recovery.

Policy implications. Our analysis shows that supply chain disruptions generate stagflation, accompanied by an increase in spare capacity for producers. This higher spare capacity curtails the supply of goods to retailers and results in a surge in prices, leading to a tighter product market. We show that, in this situation, prices become highly sensitive to changes in demand, while output remains relatively inelastic. In other words, disruptions to the supply chain enhance the effectiveness of contractionary monetary policy in taming inflation while reducing the sensitivity of output to the policy. Our results reinforce the general findings on the state-dependence of the efficacy of monetary policy (Benigno and Ricci, 2011; Liu et al., 2019; Eichenbaum et al., 2022;

Ikeda et al., 2024; Kharroubi and Smets, 2024).

We test our theoretical predictions on the enhanced effectiveness of monetary policy during supply chain disruptions by developing a threshold vector autoregression (TVAR) model that estimates the statistical differences in the effects of a contractionary monetary policy shock at different levels of the ACR index. Consistent with the theory, we find that an exogenous tightening of monetary policy leads to a significantly larger and more persistent decline in goods inflation for a given decrease in output during periods of supply chain disruptions. Our results support a more aggressive approach to tightening monetary policy in response to supply chain disturbances. As a robustness analysis, we show that we get similar results using linear projections.

Our policy results apply well beyond the COVID-19 recession: they suggest that central banks should respond vigorously to future global supply chain disruptions. In fact, our result resembles the celebrated analysis by Keynes (1940). Keynes argued that when output is constrained (in our case, because of supply chain disruptions, in Britain's case in 1940, because of resources employed in World War II), policymakers can lower aggregate demand aggressively to prevent inflation without much fear of lowering production.

Related literature. Our analysis 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 the amount of spare-labor capacity (Benigno and Eggertsson, 2023), shortages in the goods market (Blanchard and Bernanke, 2023), a quasi-kinked demand curve for produced goods (Harding et al., 2023), and capacity constraints (Comín et al., 2023). The common finding across these studies is that the scarcity of goods during disturbances to the supply chain brings the economy close to its capacity constraint, thus generating a non-linear and strong increase in inflation with a limited effect on output.

Furthermore, our paper is related to studies showing that transportation costs are important for international trade and economic activity (Allen and Arkolakis, 2014; Brancaccio et al., 2020; Dunn and Leibovici, 2023), infrastructure investment (Fuchs and Wong, 2022; Brancaccio et al., 2024), asset prices (Smirnyagin and Tsyvinski, 2022), working capital (Antràs, 2023; Kim and Shin, 2023), inflation expectations (Acharya et al., 2023), the design of new taxes and pricing rules to offset distortionary effects on the transportation network (Brancaccio et al., 2023), the interlinks between oil shocks and congestion in the supply chain (Bai and Li, 2022; Li et al.,

2022), and the effects of supply chain disruptions during the COVID-19 pandemic (Finck and Tillmann, 2022; Gordon and Clark, 2023; Finck et al., 2024).

The remainder of the paper is organized as follows. Section 2 constructs our ACR index that measures the state of the global supply chain. Section 3 develops our theoretical model and the identification restrictions for structural shocks. Section 4 presents the estimation results. Section 5 studies the state-dependent effects of a contractionary monetary policy shock following supply chain disruptions. Section 6 concludes. An extensive appendix provides further details. Our data and additional results are available on our website: https://globalportcongestion.github.io/blog/intro.html.

2. Measuring the State of the Global Supply Chain

In this section, we propose a novel index to assess the state of the global supply chain through the lens of containerized trade. More concretely, we use satellite data on the positions, speeds, and headings of container ships to measure congestion in major ports around the world.

We will start by explaining why we look at containerized trade and highlighting several key aspects of the industry. Next, we will introduce the satellite data we use. The core of the section is the motivation behind using port congestion to track the state of the global supply chain and the presentation of the algorithm that accomplishes such a goal. We will close by reporting our ACR index, discussing several aspects of our measurement, including a comparison with alternative indices in the existing literature that monitor the state of the global supply chain.

2.1. Containerized Seaborne Trade: Some Basic Facts

Containerized seaborne trade plays a prime role in the global supply chain, accounting for around 46% of all international trade (Notteboom et al., 2022).[3] In the U.S., container shipping carries more tonnage (nearly one billion short tons) and value (more than 0.7 trillion dollars) than any other means of transportation, accounting for over 50% of U.S. trade by weight and approximately 30% by value (Bureau of Transportation Statistics, 2021). Also, despite the fact that some high-value items in key shortage during the COVID-19 pandemic – e.g., computer chips – are

[3]Most of the rest is either bulk cargo (e.g., oil, grain, ore, and coal) or specialized vessels (e.g., roll-on/roll-off vessels for wheeled cargo).

often shipped by air, these items require other components to accomplish final goods – e.g., motherboards or hard drives – that are transported by container ships.

Importantly, as Brancaccio et al. (2020, p.2) explain, "[t]he 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 are built around the seaports that serve as international hubs for freight collection and distribution. Even mild congestion at these ports can impair regular supply chains and trade flows, which run under tight schedules. Any disruption leads to elevated delay costs and far-reaching trickle-down consequences for international trade and macroeconomic outcomes.

Prior to 2020, waiting times at ports were just a few hours. However, general disruptions related to the COVID-19 pandemic led to extended delays, with waiting times reaching 2-3 days at several major ports, incurring substantial financial losses. Even if a wait of 2-3 days might not seem long, an analogy is a delay on a flight arriving at an airport hub just 60 minutes late: dozens of passengers will miss their connections, generating high levels of disruption.[4] Furthermore, since nearly 80% of world trade is shipped indirectly and the average shipment stops at five additional ports before reaching the final destination (Ganapati et al., 2021), the interconnectedness in the global trade network greatly amplifies the total delays from port congestion worldwide.[5]

To frame how severe the disruptions can be, it is also important to notice that the industry is surprisingly concentrated. In 2022, there were only 5,589 container ships worldwide, of which around 500 or so belong to the larger classes in terms of size.[6] Hence, the delay of even one large ship has significant consequences for global trade. For instance, the brand new MSC Loreto carries around 24,346 TEUs (a twenty-foot equivalent unit), each with a maximum amount of cargo of $21,600$ kilograms. The MSC Loreto can load up to 240 thousand tons of cargo at full

[4]Similarly to stranded airline passengers, the buyers and sellers of goods also encountered reduced transport efficiency and heightened operational costs, demurrage and detention charges, and challenges in meeting contractual obligations and market demand. In the case of shippers and freight forwarders, the delays were compounded by surcharges like the port congestion surcharge (PCS), with fees escalating up to $1,250 per container. Given that the average value of goods in a 40-foot container (the most common container type) in 2020 was around $109,000, the PCS alone was a significant cost.

[5]To get a rough sense of the severity of the delays in shipments to the U.S., consider that it normally takes an average of 28 days to ship a container from Shanghai to Los Angeles along the Trans-Pacific route (Freightos, 2024). Typically, an average shipment stops at the Port of Ningbo-Zhoushan in China, Kaohsiung in Taiwan, Busan in South Korea, and Tokyo in Japan before arriving at the Port of Los Angeles. If we assume that waiting at each intermediate port takes 2-3 days, the total additional waiting time amounts to 8-12 days, even without considering any delays at the origin and destination ports.

[6]See https://unctad.org/rmt2022 (Accessed December 29, 2023).

capacity. A historical comparison puts this massive amount of cargo in perspective. Perhaps the most famous convoy of the Battle of the Atlantic during World War II was ONS 5, which 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 around 219 thousand tons, 10% less than the cargo capacity of the MSC Loreto.[7] Any delay in the loading and unloading operations of the MSC Loreto has ramifications for the operations of tens of thousands of different firms.[8]

The escalation of port congestion during the COVID-19 pandemic can be attributed to several factors. On the one hand, port delays were triggered by upstream and downstream issues, such as personnel mobility restrictions due to stay-at-home orders, additional quarantine measures at ports that reduced handling efficiency, trucks not arriving on time to load or unload containers due to travel controls on highways, and containers left unopened at inland factories because workers were unavailable to process the items inside them. On the other hand, a large surge in the demand for tradable goods toward the end of 2021 also contributed to elevated port congestion, as shipping companies deployed more capacity to the busiest routes and increased ship visits to catch up with demand (Notteboom et al., 2021). In other words, our index of global port congestion is a powerful measurement of the state of the global supply chain, but it encodes many different shocks. Our theoretical model will provide us with the identification assumptions to determine if an increase in port congestion was due, for example, to a demand shock or supply chain shock. But before doing so, let us describe the details of the construction of the index.

2.2. AIS Data

We use satellite data from the AIS, a tracking system mandated by the IMO. International voyaging vessels larger than 300 gross tonnage must carry a transceiver that broadcasts information about the ship (Heiland et al., 2022). Each data entry includes the IMO number, timestamp, current draft, speed, heading, and geographical coordinates.[9] The AIS processes over 2,000 reports per minute and updates information as frequently as every two seconds, offering comprehensive

[7] The official historian of the U.S. Navy, Samuel Eliot Morison, wrote: "The glorious battle of a British escort group under Commander P. W. Gretton to the Westbound convoy ONS 5 is regarded by both the Allies and the Germans as a turning point in the struggle for the North Atlantic" (Morison, 1954).

[8] In Appendix A, we provide further details on the containerized shipping industry.

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

coverage of the movements of container ships around the globe from January 2017 to September 2023.[10] The positions, speeds, and headings of ships allow us to monitor vessel movements within different port zones.

2.3. A Density-Based Spatial Clustering Algorithm

The literature on maritime economics has identified port congestion as a key measure of the state of the global supply chain (Karimi-Mamaghan et al., 2020; Bai et al., 2023; Brancaccio et al., 2024). Also, the literature has suggested measuring port congestion by estimating the likelihood that a vessel will first moor in an anchorage area within the port before docking at a berth (Talley, 2009; Talley and Ng, 2016). An anchorage is a location within a port where ships can lower anchors, while a berth is a designated spot within a port where vessels moor to load and unload cargo. If port congestion were not a concern, a ship would dock at a berth immediately upon its arrival in the port to begin loading and unloading cargo. Hence, measuring port congestion requires the identification of berth and anchorage areas, a task for which practitioners before us have largely relied on nautical charts of individual ports, making it both labor-intensive and challenging to generalize to global ports with varying internal arrangements.[11] Instead, using machine learning, we develop an iterative, multi-attribute, density-based spatial clustering algorithm that is both accurate in identifying different areas within ports and applicable to ports with different morphologies.

The algorithm identifies different port areas by focusing on the density of container ships' mooring points recorded in the AIS data. Our algorithm operates in two layers of clustering. The first layer identifies high-density areas, i.e., locations where we record many AIS observations, which are considered potential berths and anchorages. The second layer determines whether these high-density areas are berths or anchorages using our domain knowledge of vessels' behaviors in

[10] Over 99% of international container shipments are transported by container ships that exceed 500 gross tonnage. Even vessels smaller than 300 gross tonnage commonly carry an AIS transponder because of the high safety bonus it brings at a low cost (around $1000 for a basic unit). Thus, our coverage is nearly universal.

[11] Specifically, nautical charts have three major drawbacks: (i) they are static and rarely updated, so new berths may not be promptly accounted; (ii) the vast number of global ports makes it nearly impossible to manually label ship dynamics; (iii) they do not precisely identify ship movements within ports. As a result, studies using nautical charts typically focus on a single port or a limited selection of ports (Chen et al., 2016; Feng et al., 2020). Alternatively, inspecting satellite images of ports can help identify berths with fixed locations and geographical boundaries. For example, as shown in Appendix B.2, we use them to validate the estimation results for the berths in the geographically challenging Port of Ningbo-Zhoushan, located to the south of Shanghai. However, satellite images are generally incapable of identifying anchorages since their locations and boundaries are constantly changing due to weather and port conditions.

port. For example, if we observe all the headings in an orderly and close fashion, we identify the area as a berth. If, instead, the headings are more random in shape, we identify the area as an anchorage.

Figure 1 illustrates this point. The left panel shows two clusters, one in orange and one in green. Both include many AIS observations, with bows (the tips of the white signs) closely aligned (to the left in the orange cluster and to the right in the green cluster), representing the two different headings of a mooring. After we superimpose the clusters over a satellite image, we can check that this is indeed a berth. The right panel shows several clusters of the AIS data points, where the headings are random, with some of them appearing in a ring shape.[12]

(a) Headings at a Berth **(b)** Headings at an Anchorage

Figure 1: Information on Headings: Two Examples

Our algorithm is designed to address two challenges that existing clustering algorithms struggle with. First is the variability in the density of ships' mooring points *across* ports due to differences in trade volume handled, frequency of vessel visits, and geographical morphologies and boundaries. Our algorithm automatically iterates and refines its clustering parameters for each port, accommodating varied port environments. Second, thanks to its two layers, our algorithm accurately distinguishes between berth and anchorage areas *within* ports despite the high density of ships' mooring points. Importantly, our algorithm is readily adaptable to other applica-

[12] Appendix B provides details on our clustering algorithm, including the pseudo-codes and a case study involving the Port of Ningbo-Zhoushan, which illustrates the effectiveness of our methodology compared to alternative approaches in identifying berth and anchorage areas in ports with different morphologies and internal arrangements.

tions, such as port handling efficiency and waiting time, canal traffic, or stress at maritime choke points (e.g., piracy at straits). More in general, the algorithm's core mechanism – transforming domain knowledge into non-spatial attributes and using them as additional metrics between data points in an iterative clustering process – offers a versatile framework for classifying clusters of varying densities with specific labels in other contexts as well (e.g., identifying disease hotspots, urban planning, and environmental monitoring).

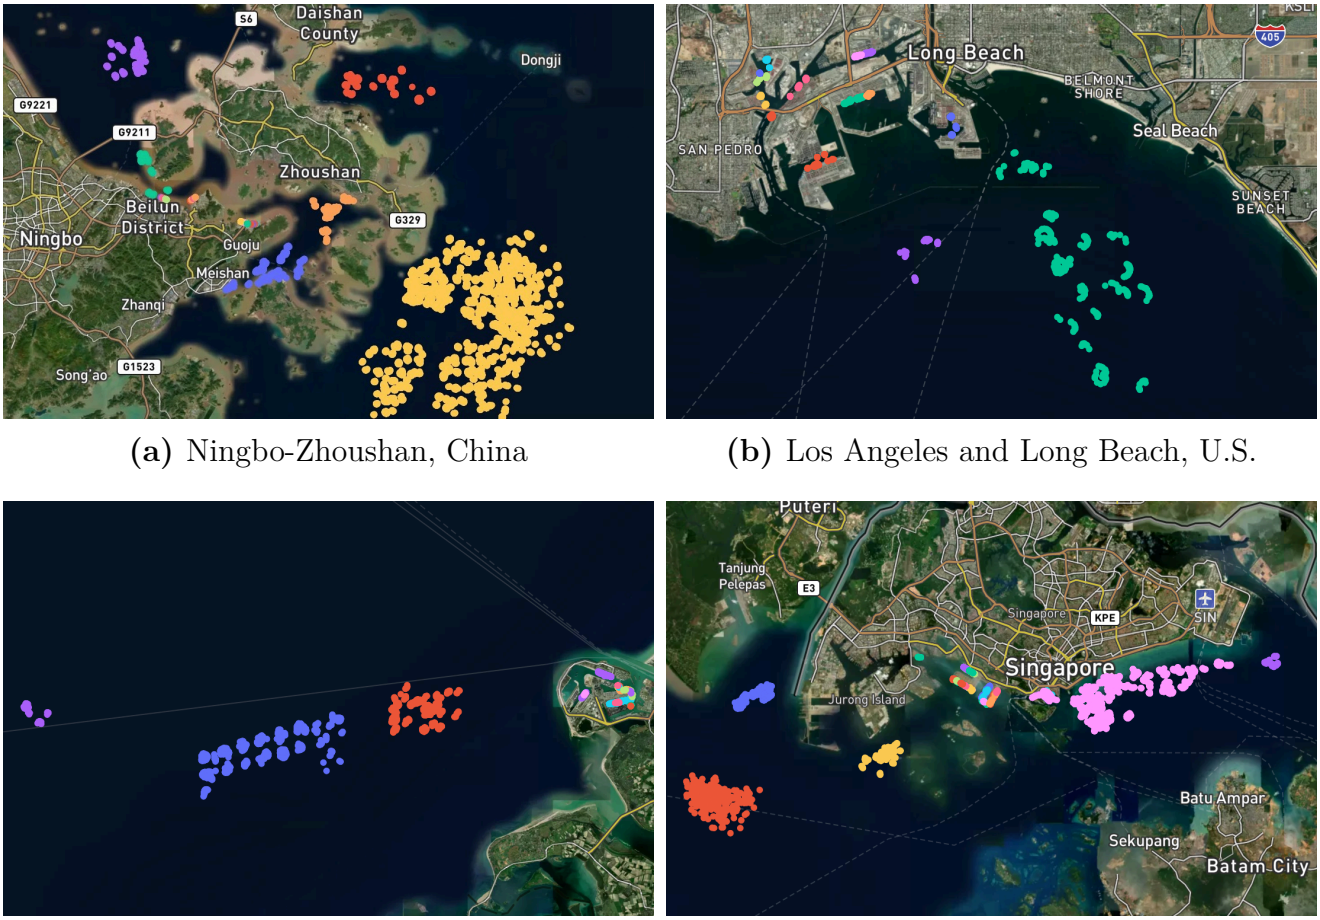

(a) Ningbo-Zhoushan, China

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

(c) Rotterdam, Netherlands

(d) Singapore

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

Note. The underlying sample for each figure incorporates the first 50,000 AIS observations of container ships entering each port since January 1, 2020.

Figure 2 reports the results of our algorithm. In each panel, we superimpose the identified anchorage (colors including 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 can also be found in Appendix B.2. Clearly, our algorithm accurately identifies the anchorage and berth areas in each port despite a broad range of geographical and operational port conditions.

2.4. The ACR Index

Port congestion arises when ships cannot immediately load and unload cargo upon arrival at ports. This delay results in vessels waiting in an anchorage area until a berth is free. For the top 50 container ports worldwide, denoted as $\mathscr{P}$, we count the number of delayed ship visits to each port p where the ship first moors in an anchorage before docking at a berth.[13] 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:

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

where $Delayed_{p,t}$ and $Undelayed_{p,t}$ represent the number of delayed and undelayed ship visits at port p in month t, respectively. We calculate the congestion rate for each port on a monthly basis throughout the sample period.

Figure 3 shows the monthly congestion rates for the top ten container ports worldwide, along with the Ports of Los Angeles, Long Beach, New York-New Jersey, and Savannah (which process more than 60% of the containerized imports to the U.S.), from January 2017 to September 2023, the period for which AIS data are available. While below we will utilize data from the top 50 ports worldwide, the ports in Figure 3 represent more than 30% of the total volume of containerized seaborne trade globally and hence summarize our main findings.

Our data indicate that the onset of the COVID-19 pandemic in March 2020 had few early effects. The congestion rates of ports such as Singapore and Rotterdam remained largely stable for several months. The situation changed substantially in the fall of 2020, as the congestion rates for the majority of ports escalated (notice the rise in chromatic intensity in Figure 3 after October 2020). In particular, our calculations indicate that approximately 80% of inbound ships at the Port of Los Angeles were unable to dock at a berth immediately upon arrival in late 2020.[14]

[13] A ship visit, or port call, refers to the arrival of a ship at a port where it docks to load and unload cargo.

[14] This measurement aligns with official statistics. According to the Pacific Merchant Shipping Association, the percentage of container ships in Los Angeles waiting five or more days for unloading surged from 10% in August to 26% in December 2020. Additionally, the Marine Exchange of Southern California reported that the number

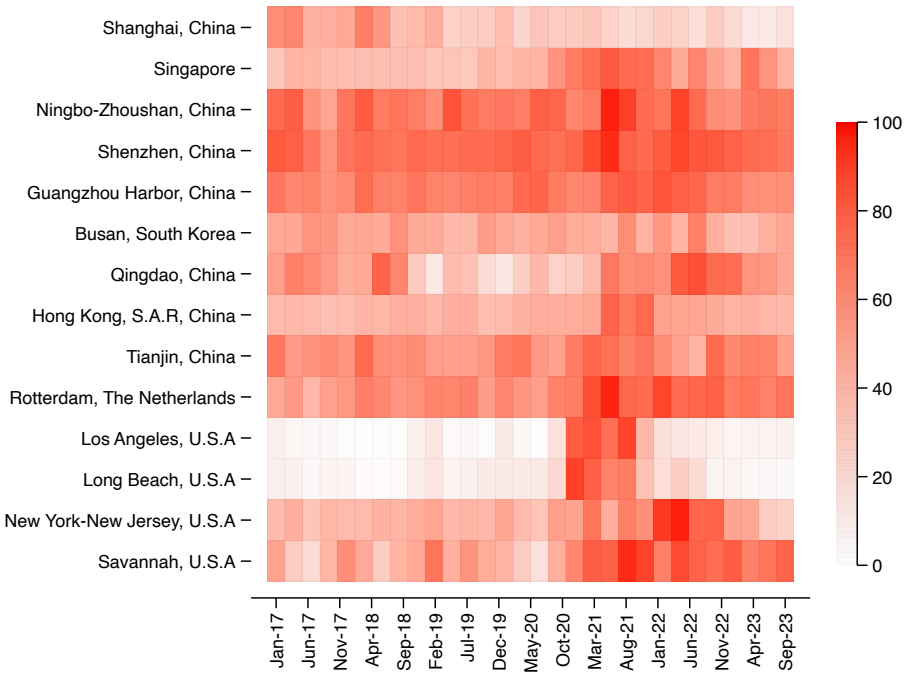

Figure 3: Congestion Rates for the Major Container Ports Worldwide

Notes. Heatmap of the 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 2017 to September 2023. The congestion rate for each port is normalized and expressed as a percentage of its peak value observed within the sample period. Cells in darker shades indicate higher congestion levels as defined in Equation (1) for the respective port during the specified month.

In sum, our analysis reveals that port congestion became acute worldwide in 2020:Q4.

To construct a time series of global port congestion that can be used in our causality analysis, we define the average congestion rate (ACR) by computing the weighted average of the congestion rates for the top 50 container ports worldwide, using as weights the relative number of ship visits to each port:

$$ACR_t = \sum_{p \in \mathscr{P}} \left[\frac{Delayed_{p,t} + Undelayed_{p,t}}{\sum_{p \in \mathscr{P}} (Delayed_{p,t} + Undelayed_{p,t})} \cdot Congestion_{p,t} \right]. \tag{2}$$

Figure 4 plots the ACR index. Prior to 2018, the index followed a declining trend, stabilizing around 28% before dropping to a sample minimum of 25% from early 2019 to mid-2020, reflecting the significant investments made worldwide in previous years to increase port capacity and supply chain resilience. Subsequently, the index consistently rose, peaking at 37% in June 2021, indicative of significant strains on the global supply chain related to the COVID-19 pandemic. At the time,

of vessels anchored in Los Angeles waters rose from fewer than 20 in August to more than 35 in December 2020.

nearly one in every three container ships that entered one of the top 50 container ports experienced delays due to port congestion. The ACR index remained elevated until mid-2022, then began to decline, returning to the sample median (29.1%) by the end of the sample period. By then, despite remaining above the average of the pre-COVID period, port congestion had returned to normal levels, and the global supply chain had resumed operating smoothly.

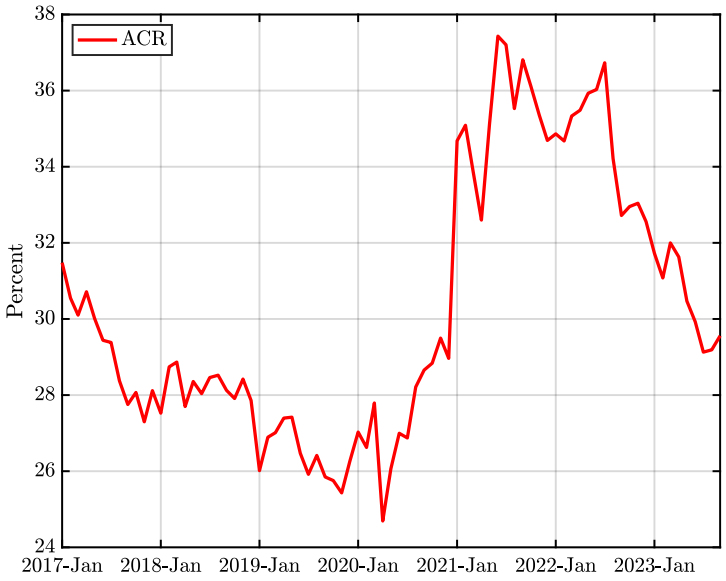

Figure 4: ACR Index

Note. The ACR index is derived by taking a weighted average of the congestion rates for the top 50 container ports worldwide, with the relative number of ship visits used as the weight for each port. The index is presented in percentage terms and has been seasonally adjusted. For the complete ranking of container ports, see https://www.worldshipping.org/top-50-ports (Accessed June 15, 2022).

2.5. Economic Consequences of Port Congestion

Port congestion has economic consequences relevant to our modeling of supply chain disruptions as heightened transportation costs or reduced matching efficiency between trading partners.

Shipping prices are significantly sensitive to port congestion given the highly concentrated market for container ships that is operated by only 5,589 vessels worldwide. Consequently, a supply chain disruption shock that raises port congestion will result in a marked shortage of shipping services and trigger a surge in shipping prices. Figure 5 shows that the Harper Peterson Time Charter Rates Index (HARPEX) – a composite indicator of container shipping rates – and our ACR index strongly co-move, with a sample correlation of 0.85 since the onset of the pandemic. This strong co-movement supports our approach of using the broad increase in transportation

costs to model supply chain disruptions.

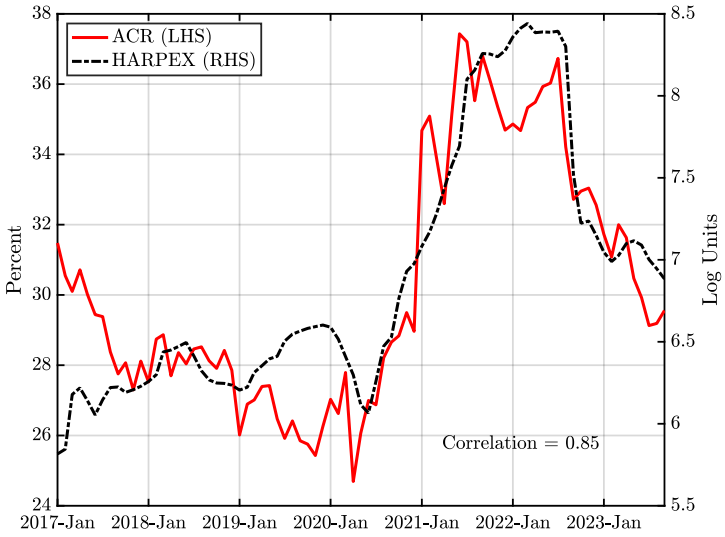

Figure 5: ACR vs. HARPEX

Notes. Figure 5 plots the ACR index (solid red line) against the HARPEX (dashed-dotted black line) for the sample period from January 2017 to September 2023. The ACR index is computed using the AIS data of container ships and our IMA-DBSCAN algorithm, as detailed in Appendix B. The HARPEX series is published by Harper Peterson and retrieved from the Refinitiv data platform. The ACR index and HARPEX are presented in percentage terms and log units, respectively. Both series have been seasonally adjusted.

Port congestion and severe delays in the delivery of production inputs also threaten the continuation of trading relationships and weaken the efficiency of establishing new partnerships. For instance, Smirnyagin and Tsyvinski (2022) and Liu et al. (2024) show that acute delivery delays resulted in the severance of an extensive number of vendor-buyer relationships in the U.S. during the COVID-19 pandemic. To resume production, buyers need to procure new suppliers while regular links are impaired, involving high searching costs for completing contractual arrangements, certifications, and regulatory compliance procedures.[15] Thus, global supply chain disruptions also manifest as a decrease in matching efficiency between trading partners. These findings motivate our alternative approach to modeling supply chain disturbances as a reduction in matching efficiency between producers and retailers.

2.6. Discussion

Several aspects of our index deserve further discussion. First, we use the relative number of ship visits to each port as weights in the calculation of the ACR index because they reflect the

[15] Xu et al. (2024) also show a heightened severance of trading relationships during supply chain disturbances.

importance of different ports within the global supply chain. A slight increase in the congestion rate for the Port of Hong Kong might signify a more pronounced global supply chain disruption than a significant increase for the Port of Manila. Nonetheless, we could build regional or national indices; our methodology would remain unchanged.

Second, the stringent contractual terms of shipping services and the fixed itineraries of container ships make our congestion tracking at seaports independent of demand and productive capacity adjustments in the *very short run*. Liner shipping companies make "tactical-level" decisions (e.g., setting the frequency of shipping services, deploying ships on their itineraries, determining sailing speeds, and designing schedules) in response to changing container shipment demand only every three to six months (Stopford, 2008; Meng et al., 2014). The actual impact on arrivals to ports might take even more months (i.e., the container ship needs to travel through its new route from Rotterdam to Los Angeles and Long Beach). In comparison, our identification assumption in Section 4 *only* requires that shipping arrivals respond with a lag of *at least one month*.

To illustrate this point, consider the Port of Shenzhen in early 2020. Despite a large contraction in world demand and a significant reduction in China's productive capacity due to strict lockdown measures and factory closures at the onset of the COVID-19 pandemic, port congestion remained high until mid-2021, as seen in Figure 3. This was because international shipping lines continued to fulfill fixed schedules and itineraries, ensuring that vessels arrived on time. However, as supply chain disruptions became severe, shipping companies adjusted routes and retailers altered procurement terms in new contracts, eventually leading to a reduction in ship visits to the Port of Shenzhen, as shown in Figure 6. Consequently, congestion in Shenzhen returned to its pre-COVID levels.

Third, the container shipping industry follows a practice known as "hurry up and wait": despite forewarnings of potential delays at the destination port, container ships often do not alter their route or speed, as doing so would necessitate changing many contractual arrangements. Hence, even if ports started to become congested in the fall of 2020, it is safe for our empirical analysis to consider the routes and speeds of vessels as fixed for at least one month. Furthermore, Appendix A documents the unimportance of oil prices for ship speeds, corroborating a well-known result in the literature.

Lastly, since the draft of a vessel normally reflects its cargo load, we could potentially augment

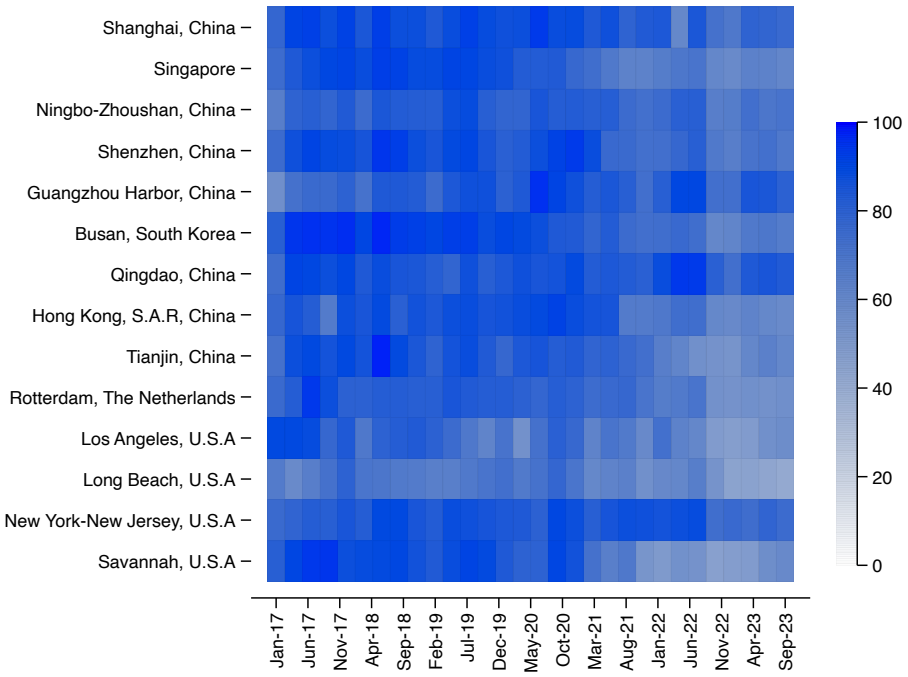

Figure 6: Ship Visits to the Major Container Ports Worldwide

Note. The heatmap displays the number of ship visits to each of the top ten global container ports, as well as the Ports of Los Angeles, Long Beach, New York-New Jersey, and Savannah, covering the period from January 2017 to September 2023. The number of ship visits to each port is normalized and expressed as a percentage of its peak value observed within the sample period. Cells in darker shades denote a greater number of ship visits to the respective port during the specified month.

the current congestion indices with this information and, hence, compute the weight of delayed cargo at major ports worldwide. Yet, unlike dry bulk ships or oil tankers (Bai and Li, 2022; Li et al., 2022), which only take on new shipping orders once they have finished the previous ones ("taxis" as referred to in Brancaccio et al., 2020), implying large fluctuations in a vessel's draft when it loads *or* unloads cargo, container ships load *and* unload cargo simultaneously, often resulting in small changes in their draft at port. Consequently, the weight-augmented congestion indices would behave largely the same as the original congestion indices. Appendix A discusses additional issues in more detail, such as the rate of ship idleness.

2.7. Alternative Indices

Appendix F compares our ACR index to other popular indices in the literature that track the state of the global supply chain, notably the HARPEX, New York Fed's Global Supply Chain Pressure Index (GSCPI), and the Supply Disruptions Index (SDI) compiled by Smirnyagin and

Tsyvinski (2022) and Liu et al. (2024). We show that there are significant differences between the indices that influence the interpretation of supply chain disruptions and their causal effects on macro aggregates, including inflation. Additionally, Appendix F.5 develops a targeted ACR index for the major ports along the Trans-Pacific route, which is the most important shipping route connecting East Asia (mostly China) and the U.S. Appendix F.6 also constructs an alternative measure of port congestion – the average congestion time (ACT) – using AIS data and our spatial clustering algorithm. Consistent with the literature (Brancaccio et al., 2024), the ACT index measures the average number of hours a container ship waits in the anchorage area of a port before docking at a berth, weighted by the relative number of ship visits to each port. We show that using either the targeted ACR index or the ACT index in the causality assessment delivers results quantitatively similar to those obtained with our original ACR index.

The integration of high-frequency AIS data with our spatial clustering algorithm also enables the construction of port congestion indices at even higher frequencies than monthly updates. Appendix B.3 reconstructs the ACR and ACT indices using weekly updates of the AIS data and highlights that, despite heightened volatility, the weekly indices of port congestion exhibit the same patterns as their monthly counterparts.

3. A Model of Congestion and Spare Capacity

Next, we develop a model of congestion and spare capacity that will provide us with identification restrictions for our causality analysis. Without theoretical guidance, it is hard to find a research design that allows us to ascertain causality for the shocks to aggregate demand, supply, and the supply chain simultaneously, given that their impacts on macro aggregates are often intertwined, particularly during the COVID-19 pandemic.

We develop a search and matching model because it highlights what was, to us, the fundamental issue during the COVID-19 pandemic: the inability of prices to play a fully allocative role. Nonetheless, our model will impose identification restrictions that could also be derived from a New Keynesian model. In such a framework, nominal rigidities replace search and matching frictions as impediments to the smooth working of prices. Readers who prefer the New Keynesian thinking shall not find it too difficult to jump to Table 1, where we summarize the identification restrictions derived from our theory, and see how those restrictions would also come from many

New Keynesian models.

Our economy consists of producers, retailers, and households, with producers based in a different location from both retailers and households. Producers manufacture goods using a fixed-factor endowment and incur transportation costs when selling the goods to retailers. Retailers purchase goods from producers but face search frictions that make it difficult to meet with the former. Retailers then sell the goods to households. Households own the producers and retailers, accruing all their profits earned through trades.

We distinguish between producers and retailers in different locations to capture the idea that firms must trade in a global supply chain. Search and matching frictions make this trade non-trivial.[16] Transportation costs and search frictions hinder the allocative role of prices in clearing the quantity of goods sold by producers to retailers. Thus, we introduce disruptions to the global supply chain either as higher transportation costs or lower matching efficiency between producers and retailers.[17]

Our model will allow us to study three standard shocks: an aggregate demand shock (i.e., a change in the money held by households due to a change in monetary policy or a change in their preference for consuming goods), a productive capacity shock (i.e., a change in the fixed-factor endowment of producers), and a supply chain shock (i.e., an increase in transportation costs or lower matching efficiency). Our model has distinctive predictions regarding the co-movements of spare productive capacity with prices and output in response to each of these shocks that will let us identify the causal effects of supply chain disruptions (Section 4) and connect these disruptions to the effectiveness of monetary policy (Section 5).

3.1. Producers and Retailers

There is an exogenous unit mass of producers and an endogenous measure of retailers. When matched with a retailer, a producer manufactures $y = l$ final goods using its fixed-factor endowment.[18] Producers sell the goods to retailers in a frictional goods market that prevents the sale

[16] Our approach is similar to the case where firms require intermediate goods for the production of final goods. See, for instance, Costinot et al. (2013), Kasahara and Lapham (2013), and Ramondo and Rodríguez-Clare (2013).

[17] These modeling choices are also consistent with our discussion in Section 2.5 regarding the economic manifestations of escalated port congestion. Appendix C also discusses the evidence of search and matching frictions in the goods market and the relevance of transportation costs for the severance of commercial trade.

[18] We abstract away from modeling the endogenous production decisions of producers, as this would require a multi-country, multi-sector model of production networks to delve into the propagation of sectoral supply disruptions across countries (Baqaee and Farhi, 2022; di Giovanni et al., 2022, 2023), which is informative but not the

of the 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 $\rho > 0$. Upon a meeting between a producer and a retailer that results in a trade (we will discuss below when this happens), the retailer resells the purchased good to the household at a price p.

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

$$\mathscr{M} = A(x_U^{-\xi} + i_U^{-\xi})^{-\frac{1}{\xi}}, \tag{3}$$

where A is the efficiency of the matching function and ξ is the elasticity of substitution between x_U and i_U. We assume $\xi > 0$, such that $\mathscr{M} \leq A \cdot \min\{x_U, i_U\}$. For the moment, we parameterize $A = 1$. Later, we will explore a supply chain shock that yields $0 < A < 1$.

Product market tightness $\theta \equiv i_U/x_U$ is the ratio between the number of visits by the unmatched retailers and the number of unmatched producers. Product market tightness is taken as given by individual firms. Specifically, the probability for a producer to meet a retailer is:

$$f(\theta) = \frac{\mathscr{M}}{x_U} = A(1 + \theta^{-\xi})^{-\frac{1}{\xi}}, \tag{4}$$

and the probability for a retailer to meet a producer is:

$$q(\theta) = \frac{\mathscr{M}}{i_U} = A(1 + \theta^{\xi})^{-\frac{1}{\xi}}. \tag{5}$$

The function $f(\theta)$ satisfies that $f(0) = 0$, $\lim_{\theta\to+\infty} f(\theta) = A$, and $f'(\theta) > 0$, whereas $q(\theta)$ satisfies that $q(0) = A$, $\lim_{\theta\to+\infty} q(\theta) = 0$, and $q'(\theta) < 0$. Two additional properties that will be useful later are that $f(\theta)/q(\theta) = \theta$ and $f'(\theta) = A^{-\xi}q(\theta)^{1+\xi}$.

Transportation cost. Producers pay a per-unit idiosyncratic transportation cost to ship their goods to retailers.[19] In each period, producers draw a per-unit transportation cost z from the log-normal distribution $G(z)$ with the scale parameter γ and the shape parameter σ, i.e.,

focus of our model. Instead, our emphasis is on the severance of trading relationships between producers and retailers in different locations – due to high transportation costs or reduced matching efficiency – that leads to global supply chain disruptions. Such an exogenous production approach preserves the core predictions regarding the impact of an aggregate supply shock on selected macro aggregates, which can then be directly used as identification restrictions in the SVAR model, thereby simplifying the complexities associated with modeling production networks.

[19] Our results hold if the transportation cost is borne by retailers instead because the match separation condition (14) is invariant to this modeling choice. For simplicity, we also assume that the household receives this shipping cost as a payment for its work in moving the goods.

$G(z) \equiv \Phi\left[(\log z - \gamma)/\sigma\right]$, where $\Phi(\cdot)$ is the standard normal cumulative density function.[20] As we discuss later, there exists a reservation level of transportation cost $\bar{z}$, above which matches are unprofitable and, hence, severed ($z > \bar{z}$), whereas they continue otherwise ($z \leq \bar{z}$).

Value functions. At the beginning of each period, the matched producers sell the manufactured goods to retailers and pay the transportation costs. In contrast, the matched retailers sell their purchased goods to households and pay the wholesale price of goods to the producers. The unmatched producers and retailers search to form a match with each other. At the beginning of the next period, each producer draws a new transportation cost, and the match continues if the new cost is sufficiently low, such that there is a positive surplus from trade.[21]

Four value functions describe the return for the different statuses of producers and retailers. The value for a matched producer (identified by the subscript M), $X_M(z)$, is equal to:

$$X_M(z) = (r(z) - z)\, l + \beta \mathbb{E}_{z'} \left[\max\left(X_M(z'), X_U\right)\right], \tag{6}$$

where $r(z)$ is the (endogenous) wholesale price per unit of the final good, β is the discount factor, and z' is the draw of transportation cost at the beginning of the next period. Equation (6) shows that the present value of being a matched producer is the profit margin $(r(z) - z)\, l$, plus the continuation value, which depends on whether the producer separates from the match. Separation is determined by the transportation cost next period, z', and the max operator picks the optimal continuation/separation decision.

The value for an unmatched producer, X_U, is:

$$X_U = \beta f(\theta) \mathbb{E}_{z'} \left[\max\left(X_M(z'), X_U\right)\right] + \beta\left(1 - f(\theta)\right) X_U. \tag{7}$$

[20] The modeling assumption that transportation costs follow a log-normal distribution is consistent with Kasahara and Lapham (2013). A random distribution of transportation costs captures the idea that, in the real world, there is a wide range of transportation costs that vary across countries, routes, directions, and commodities (Brown et al., 2021). This approach is more appealing than a fixed transportation cost (such as the "iceberg" formulation of trade costs in Samuelson, 1954). Additionally, a log-normal distribution includes a scale parameter such that an increase in this parameter raises transportation costs in general, thus serving as a possible theoretical counterpart to our ACR index. This setting allows us to model an exogenous increase in transportation costs conveniently and, more importantly, preserves the positive empirical relationship between the ACR index and shipping costs, as highlighted in Figure 5. Alternatively, we could augment the current model with a full-fledged transportation sector, where interactions between producers, shipowners, and retailers determine the transportation cost. Such an endogenous setting can be found in Brancaccio et al. (2020), Bai and Li (2022), and Dunn and Leibovici (2023). Nonetheless, we maintain the current setting for its tractability.

[21] We could consider a more general setup where each producer maintains its previous draw of transportation cost with a given probability and the producer draws a new transportation cost from $G(z)$ with the complementary probability. This setup is often found in models that study the labor market outcomes following a rise in economic turbulence (den Haan et al., 2005; Thomas and Zanetti, 2009; Fujita, 2018; Pizzinelli et al., 2020). Despite more involved algebra, our main results still hold in this more general setup.

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

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

$$I_M(z) = (p - r(z))\, l + \beta \mathbb{E}_{z'} \left[\max \left(I_M(z'), I_U \right) \right]. \tag{8}$$

The retailer earns the price p by reselling each unit of the purchased goods to the households and pays the corresponding wholesale price $r(z)$ to the producer. As before, the max operator picks the optimal continuation/separation decision conditional on z'.

If the drawn transportation cost makes the match unprofitable, the retailer separates from the match and starts the next period with a return:

$$I_U = -\rho l + \beta q(\theta) \mathbb{E}_{z'} \left[\max \left(I_M(z'), I_U \right) \right] + \beta \left(1 - q(\theta) \right) I_U, \tag{9}$$

where ρ is a fixed cost per unit of the final good that the retailer pays to the producer during the visit. Free entry into the product market drives the value for an unmatched retailer to zero in equilibrium, i.e., $I_U = 0$.

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

$$S(z) = X_M(z) - X_U + I_M(z) - I_U, \tag{10}$$

and it is split through Nash bargaining. The producer earns a constant share η of the total surplus, and the retailer earns the remaining share $1 - \eta$, which in equilibrium yields:

$$\eta \left(I_M(z) - I_U \right) = (1 - \eta) \left(X_M(z) - X_U \right). \tag{11}$$

Given the Nash bargaining sharing rule (11), the value functions (6), (7), (8), and the free-entry condition $I_U = 0$, the wholesale price that splits the surplus is equal to:

$$r(z) = \eta(p + \rho\theta) + (1 - \eta)z. \tag{12}$$

When the bargaining power of the producer is low ($\eta \to 0$), the wholesale price is close to the cost of transportation (z). Congestion in the matching process, captured by tightness in the product

market, worsens the bargaining position of retailers by lowering their matching probability. Thus, higher tightness increases the wholesale price retailers pay to the producers.

Match separation. Since the total value for a matched producer and a matched retailer, i.e., $X_M(z) + I_M(z)$, strictly decreases with the cost of transportation z, there exists a cut-off transportation cost $\bar{z}$, above which the costs are too high, making the matches unprofitable and consequently severed. This cut-off makes the total surplus in Equation (10) equal to zero:

$$S(\bar{z}) = 0. \tag{13}$$

By substituting the value functions (6), (7), (8), and the free-entry condition $I_U = 0$ into Equation (13), we can express the match separation condition as a function of p, $\bar{z}$, and θ satisfying:

$$\mathbb{F}(p, \bar{z}, \theta) = (p - \bar{z})\, l + (1 - \eta f(\theta))\, \beta \mathbb{E}_{z'} S(z') = 0, \tag{14}$$

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 (9) and the free-entry condition $I_U = 0$, we define the match creation condition as a function of $\bar{z}$ and θ satisfying:

$$\mathbb{H}(\bar{z}, \theta) = \frac{\rho l}{q(\theta)} - (1 - \eta)\beta \mathbb{E}_{z'} S(z') = 0. \tag{15}$$

Aggregate supply. The aggregate supply in the economy results from the equilibrium in the product market, defined as:

Definition 1. *The equilibrium in the product market consists of a price p, a reservation transportation cost $\bar{z}$, and a product market tightness θ such that the conditions for match separation (14) and match creation (15) simultaneously hold: $\mathbb{F}(p, \bar{z}, \theta) = \mathbb{H}(\bar{z}, \theta) = 0$.*

Definition 1 tells us that the equilibrium product market tightness can be expressed as a function of the price and the reservation transportation cost. Next, we characterize this relationship.

Proposition 1. *In equilibrium, the price p, reservation transportation cost $\bar{z}$, and product market tightness θ satisfy:*

$$\theta(p, \bar{z}) = \frac{1 - \eta}{\eta \rho} \left(p - \bar{z} + \beta \int_0^{\bar{z}} G(z') dz' \right), \tag{16}$$

where $G(\cdot)$ is the log-normal cumulative density function. Hence, the equilibrium product market tightness θ has the following properties with respect to the price p:

1. *$\theta(p^{min}, \bar{z}) = 0$ and $\lim_{p\to+\infty} \theta(p, \bar{z}) = +\infty$, where p^{min} satisfies:*

$$p^{min} - \bar{z} + \beta \int_0^{\bar{z}} G(z')dz' = 0;$$

2. *$\theta(p, \bar{z})$ is strictly increasing on $[p^{min}, +\infty)$; and*

3. *$\theta(p, \bar{z})$ is linear on $[p^{min}, +\infty)$.*

Proof. See Appendix D.1. ■

Proposition 1 establishes that the equilibrium product market tightness strictly increases with the price of goods. This is intuitive. When the total surplus rises due to a higher price, retailers visit more producers, leading to increased tightness. Additionally, in Proposition 1′ reported in Appendix D.8, we show that the equilibrium product market tightness decreases with the reservation transportation cost. This is because a rise in the reservation transportation cost shrinks the total surplus shared between producers and retailers at the margin. As a result, the incentives for retailers to visit producers diminish, causing a slacker product market.

Next, the aggregate supply comprises the quantity of goods traded by the retailers and producers that survive separation for a given productive capacity, equal to the total factor endowment of producers l. To determine the equilibrium number of matched producers, we consider the law of motion for the number of matched producers at the beginning of the next period:

$$x'_M = G(\bar{z})x_M + f(\theta)G(\bar{z})x_U, \tag{17}$$

and that for the number of unmatched producers at the beginning of the next period:

$$x'_U = \left[1 - f(\theta) + f(\theta)\left(1 - G(\bar{z})\right)\right] x_U + \left(1 - G(\bar{z})\right) x_M. \tag{18}$$

Instead of examining the full transition dynamics of the model (a discussion we relegate to Appendix D.10), we focus on the steady state. As shown in Section 3.4, using comparative statics suffices to derive a set of identification restrictions for each shock of interest in our causality analysis.

Setting $x'_M = x_M$ in Equation (17) and recalling that $x_M + x_U = 1$, we derive the steady state number of matched producers:

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

where the product market tightness θ, as determined in Equation (16), will be solved as a (steady state) constant once we impose the goods market clearing condition. The (steady state) aggregate supply is thus equal to the quantity of goods supplied by the (steady state) number of matched producers given l:

$$c_s(\bar{z},\theta) = x_M^{ss}(\bar{z},\theta)l = \frac{f(\theta)G(\bar{z})}{1-G(\bar{z})+f(\theta)G(\bar{z})}l. \tag{19}$$

By substituting the expressions for $f(\theta)$ and θ from Equations (4) and (16) into Equation (19), we express the aggregate supply as a function of price and reservation transportation cost.

Definition 2. *The aggregate supply c_s, expressed as a function of p and $\bar{z}$, is equal to:*

$$c_s(p,\bar{z}) = \frac{A\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\bar{z}+\beta\int_0^{\bar{z}}G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}G(\bar{z})}{1-G(\bar{z})+A\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\bar{z}+\beta\int_0^{\bar{z}}G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}G(\bar{z})}l, \tag{20}$$

for all $(p,\bar{z})$ satisfying:

$$p-\bar{z}+\beta\int_0^{\bar{z}}G(z')dz' \geq 0. \tag{21}$$

Since $c_s(p,\bar{z})$ is determined by two endogenous variables, there exist infinite combinations of p and $\bar{z}$'s that yield the same aggregate supply, as long as they satisfy the constraint (21). The reason is that each producer-retailer pair decides the price in a situation of bilateral monopoly, a problem with indeterminate solution (Howitt and McAfee, 1987; Hall, 2005). This indeterminacy is common in related search models (e.g., Michaillat and Saez 2015).

We resolve the indeterminacy by selecting the equilibrium (and its associated steady state) where the reservation transportation cost remains fixed at an arbitrary level τ, and the price moves to satisfy the aggregate supply condition.[22] Selecting one equilibrium (at least implicitly) by determining one variable from outside the model is standard in the search literature. For example, in Mortensen and Pissarides (1994), the price of the final good is assumed to follow an exogenous stochastic process.[23]

By considering the equilibrium (and its associated steady state) with freely adjusting prices, we can study the responses of prices to the distinct disturbances to aggregate demand, productive capacity, and the supply chain, respectively, and then use their co-movements with other

[22] In Appendix C, we provide an economic interpretation of the exogenous reservation transportation cost. In short, it relates to the minimum transportation expense necessary for engaging in trade, including costs such as fuel, labor, insurance, and tariffs, which are largely fixed in the short term.

[23] Another possibility is to pick the price of the final good as the numeraire in the economy. We do not follow this route here because we want to build a theory of aggregate demand where the monetary unit is the numeraire.

endogenous variables to formulate unique identification restrictions to estimate the causal effects of supply chain disturbances in our SVAR model in Section 4, as well as the state-dependent effects of monetary tightening in our TVAR model in Section 5. Hence, Definition 2′ recasts the original Definition 2 of the aggregate supply as a function of price p for an arbitrary reservation transportation cost τ.[24]

Definition 2′. *For an arbitrary reservation transportation cost $\tau \in (0, +\infty)$, the flexible price aggregate supply c_s^{flex} is the function of price p defined by:*

$$c_s^{flex}(p) = \frac{A\left\{1 + \left[\frac{1-\eta}{\eta\rho}\left(p - \tau + \beta\int_0^\tau G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}} G(\tau)}{1 - G(\tau) + A\left\{1 + \left[\frac{1-\eta}{\eta\rho}\left(p - \tau + \beta\int_0^\tau G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}} G(\tau)} l, \tag{22}$$

for all $p \in [p^{min}, +\infty)$, where p^{min} satisfies:

$$p^{min} - \tau + \beta\int_0^\tau G(z')dz' = 0.$$

The next proposition outlines the properties of the aggregate supply when the price adjusts to satisfy the aggregate supply condition.

Proposition 2. *The flexible price aggregate supply c_s^{flex} has the following properties:*

1. *$c_s^{flex}(p^{min}) = 0$ and $\lim_{p\to+\infty} c_s^{flex}(p) = G(\tau)l$;*
2. *$c_s^{flex}(p)$ is strictly increasing in p on $[p^{min}, +\infty)$; and*
3. *$c_s^{flex}(p)$ is concave on $[p^{min}, +\infty)$.*

Proof. See Appendix D.2. ■

The aggregate supply $c_s^{flex}(p)$ in Equation (22) represents the quantity of goods traded that satisfies Equation (20) for a given reservation transportation cost τ. That is, the interaction between the price and tightness in the product market determines the aggregate supply. A higher price leads to a greater total surplus by increasing the value for the matched retailer. This, in turn, enhances the incentives for retailers to visit producers, resulting in increased product market tightness and a higher probability for a producer to match with a retailer. More matches are

[24] Appendix D.8 discusses the alternative equilibrium selection mechanism in which p remains fixed while $\bar{z}$ can vary. In addition to deriving its key analytical properties, we use numerical methods to approximate this fixed price aggregate supply and illustrate its properties across different values of the reservation transportation cost.

created, resulting in a higher aggregate supply. While transportation costs and matching frictions reduce the aggregate supply of goods to retailers and create spare capacity, the model retains the standard positive relationship between the price and the aggregate supply.

Figure 7 shows the aggregate supply in the quantity-price (Q,p) plane. For a given productive capacity of the economy l (brown line), transportation costs limit the production to $G(\tau)l$ (green line).[25] Search frictions further reduce the aggregate supply to the level $c_s^{flex}(p)$ (blue line), which, as in standard models, increases with the price. The spare capacity, represented by the difference between the productive capacity of the economy and the actual production (i.e., $l - c_s^{flex}(p)$), arises from the presence of both transportation and matching costs.

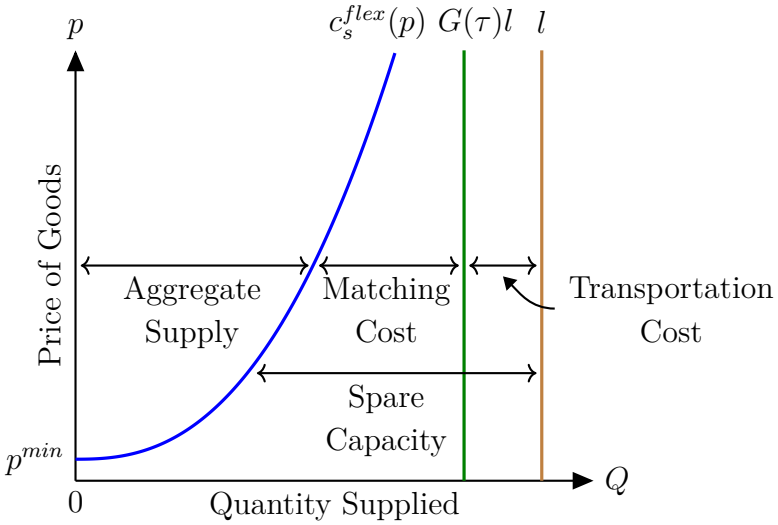

Figure 7: Supply Side of the Economy with Flexible Prices

3.2. The Representative Household

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

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

where c denotes consumption, m is the nominal money balance, p is the price level, the parameter $\chi > 0$ represents the taste for consumption relative to holding money, and the parameter $\varepsilon > 1$ is the elasticity of substitution between consumption and real money balances. We borrow this

[25] Setting $A = 1$ in Equation (22) and assuming the absence of search frictions yield such a boundary.

utility function from Michaillat and Saez (2015) to ensure that aggregate demand is instrumental to the changes in macro aggregates.

Taking the price as given, the household chooses consumption and nominal money balances to maximize utility, subject to the budget constraint:

$$pc + m \leq \mu + \underbrace{pc_s^{flex}(p) - \int_0^\tau z' c_s^{flex}(p) dG(z')}_{\text{Profits of Producers \& Retailers}} + \underbrace{\int_0^\tau z' c_s^{flex}(p) dG(z')}_{\text{Transportation Costs}} \tag{23}$$

$$= \mu + p\left[\frac{f\left(\theta(p)\right) G(\tau)}{1 - G(\tau) + f\left(\theta(p)\right) G(\tau)} l\right],$$

where $\mu > 0$ is the household's endowment of nominal money.

Solving the household's optimization problem yields the optimal condition:

$$\frac{\chi}{1+\chi} c^{-\frac{1}{\varepsilon}} = \frac{1}{1+\chi}\left(\frac{m}{p}\right)^{-\frac{1}{\varepsilon}}. \tag{24}$$

Aggregate demand. The aggregate demand in the economy is equal to the level of consumption that maximizes utility at a given price when the money market clears (this condition holds in and outside the steady state). By replacing m with μ in Equation (24) and rearranging, we derive the aggregate demand in the economy.

Definition 3. *The aggregate demand c_d for a given price $p \in (0, +\infty)$ equals:*

$$c_d(p) = \chi^\varepsilon \frac{\mu}{p}. \tag{25}$$

Proposition 3. *$c_d(p)$ is strictly decreasing and convex on $(0, +\infty)$.*

Proof. Direct proof from Equation (25). ■

Figure 8 below shows the aggregate demand, which is downward sloping in the (Q, p) plane. Since a higher price leads to lower real money balances, the household's indifference between consumption and holding money implies that it desires lower consumption when the price is higher. Hence, the aggregate demand in the economy decreases with the price.

3.3. The Flexible Price Steady State

For a given reservation transportation cost τ, the flexible price steady state is presented in Definition 4, and its existence is demonstrated in Proposition 4.

Definition 4. *Fixing the reservation transportation cost $\bar{z}$ to an arbitrary value $\tau > 0$, the flexible price steady state consists of a price p that equates aggregate supply and aggregate demand, $c_s^{flex}(p) = c_d(p)$, yielding:*

$$\frac{f\left(\theta(p)\right) G(\tau)}{1 - G(\tau) + f\left(\theta(p)\right) G(\tau)} l = \chi^{\varepsilon} \frac{\mu}{p}, \tag{26}$$

where the product market tightness θ is given by:

$$\theta(p) = \frac{1-\eta}{\eta\rho} \left(p - \tau + \beta \int_0^{\tau} G(z')dz' \right). \tag{27}$$

In addition, the household's budget constraint (23) also holds with equality.

Proposition 4. *For any $\tau > 0$, there exists a unique flexible price steady state that features positive price and consumption.*

Proof. See Appendix D.3. ■

Figure 8 shows the aggregate supply, aggregate demand, and the steady state price, p_0, where the aggregate supply and demand intersect in the (Q, p) plane. The maximum quantity of goods that could be supplied when the matching process between producers and retailers becomes frictionless and the productive capacity of the economy are also plotted in the figure for comparison.

3.4. Comparative Statics

We use comparative statics to study the responses of the macro aggregates to (unanticipated) adverse shocks to aggregate demand, productive capacity, and the supply chain, respectively, when the economy is at the steady state. These responses provide unique identification restrictions for studying the causal effects of supply chain disturbances in the SVAR model in the subsequent section, as well as the state-dependent effects of monetary tightening in the TVAR model in Section 5. Appendix D.10 shows (numerically) the complete dynamics of the model after each shock. Suffice it to say here that the transition dynamics are fully consistent with our discussion below and the implied identification restrictions.

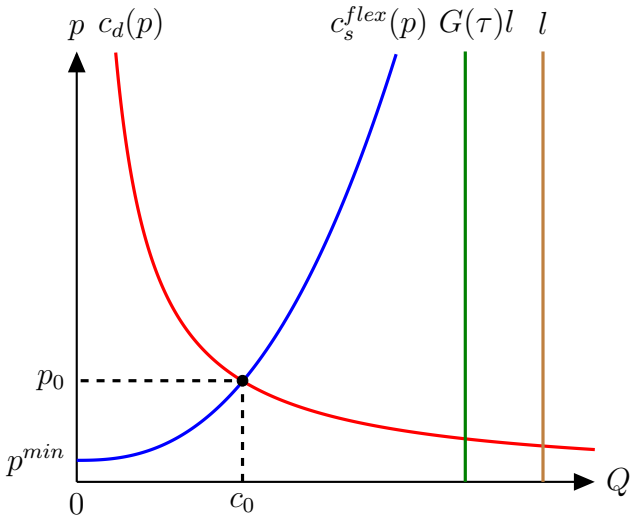

Figure 8: Aggregate Demand and Supply

An adverse shock to aggregate demand can manifest as either a decrease in the money supply, μ, or a decline in the preference for consuming goods, χ. An adverse shock to productive capacity corresponds to a negative disturbance to the fixed-factor endowment of producers, l. An adverse shock to the supply chain involves an increase in the distribution of transportation costs (characterized by a rise in γ, the scale parameter of the log-normal distribution of transportation costs) or lower matching efficiency (characterized by a fall in A such that $0 < A < 1$). Proposition 5 summarizes the responses of the macro aggregates to each shock of interest.

Proposition 5. *At the steady state:*

- ***An adverse shock to aggregate demand*** ($\mu \downarrow$ *or* $\chi \downarrow$) *increases matching cost and spare capacity while it decreases consumption (or, equivalently, output), price, product market tightness, and wholesale price.*

- ***An adverse shock to productive capacity*** ($l \downarrow$) *increases price, product market tightness, and wholesale price. At the same time, it decreases consumption (or, equivalently, output), matching cost, and spare capacity.*

- ***An adverse shock to the supply chain*** ($\gamma \uparrow$ *or* $A \downarrow$) *increases price and spare capacity while it decreases consumption (or, equivalently, output). The responses of product market tightness, wholesale price, and matching cost are undetermined after an increase in transportation costs* ($\gamma \uparrow$). *In comparison, product market tightness and wholesale price increase while matching cost remains undetermined after a fall in matching efficiency* ($A \downarrow$).

Proof. See Appendices D.4 and D.9. ■

Table 1 summarizes the signs of the responses of the endogenous variables to each shock. Figure 9 plots the comparative statics (left panels) for the aggregate demand, productive capacity, and supply chain shocks alongside the corresponding equilibrium conditions between product market tightness and price (right panels) from Equation (27), which describes the optimal response of product market tightness to a change in price.[26]

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

Adverse Shock To:	**Effects On:**					
	Consumption (or Output)	Price	Product Market Tightness	Wholesale Price	Matching Cost	Spare Capacity
	c	p	θ	r	$\frac{AG(\tau)}{1-(1-A)G(\tau)}l-c$	$l-c$
Aggregate Demand ($\mu\downarrow$ or $\chi\downarrow$)	$-$	$-$	$-$	$-$	$+$	$+$
Productive Capacity ($l\downarrow$)	$-$	$+$	$+$	$+$	$-$	$-$
Supply Chain ($\gamma\uparrow$)	$-$	$+$	Undetermined	Undetermined	Undetermined	$+$
Supply Chain ($A\downarrow$)	$-$	$+$	$+$	$+$	Undetermined	$+$

Panel (a) in Figure 9 shows the comparative statics for a decline in aggregate demand. The aggregate demand curve shifts inward from c_d to c_d', driven by the preference for lower consumption by households, either because they hold less money or prefer to decrease consumption.[27] Thus, the price decreases to clear the market. As the price decreases and the profits from sales to the households fall, retailers visit fewer producers to participate in trade, hence lowering product market tightness. The declines in price and product market tightness lead to a lower wholesale price, since not only is the sale of goods less profitable, but also the probability of establishing a match with producers increases. Consequently, producers sell a lower fraction of their productive capacity, resulting in lower consumption (or, equivalently, output) as well as higher spare capacity and matching cost.

Panel (b) in Figure 9 shows the response to a negative supply shock that decreases productive capacity from l to l'. This shock causes the aggregate supply curve to rotate inward while leaving p^{min} unchanged (since the distribution of transportation costs remains the same, and thus the minimum price for profitable transactions is unchanged). The price increases to clear the market

[26] To save space, we relegate the discussion of the case with lower matching efficiency to Appendix D.9, since the reasoning closely aligns with the scenario involving higher transportation costs.

[27] Notice that the reasoning in the main text holds also for fiscal shocks (if we incorporated a government to our model). Thus, our demand shocks encompass fiscal and monetary policy shocks as well as preference shocks.

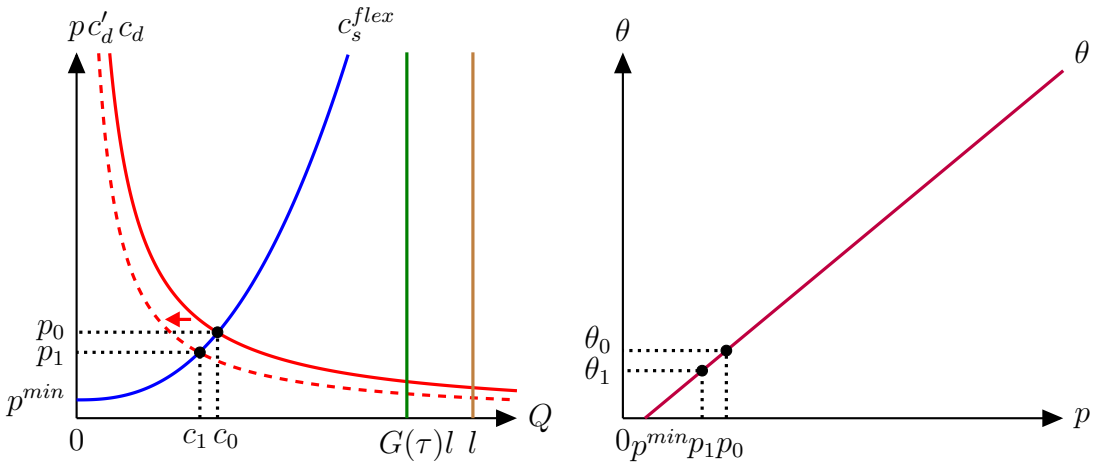

(a) An Adverse Shock to Aggregate Demand, i.e., $\mu \downarrow$ or $\chi \downarrow$

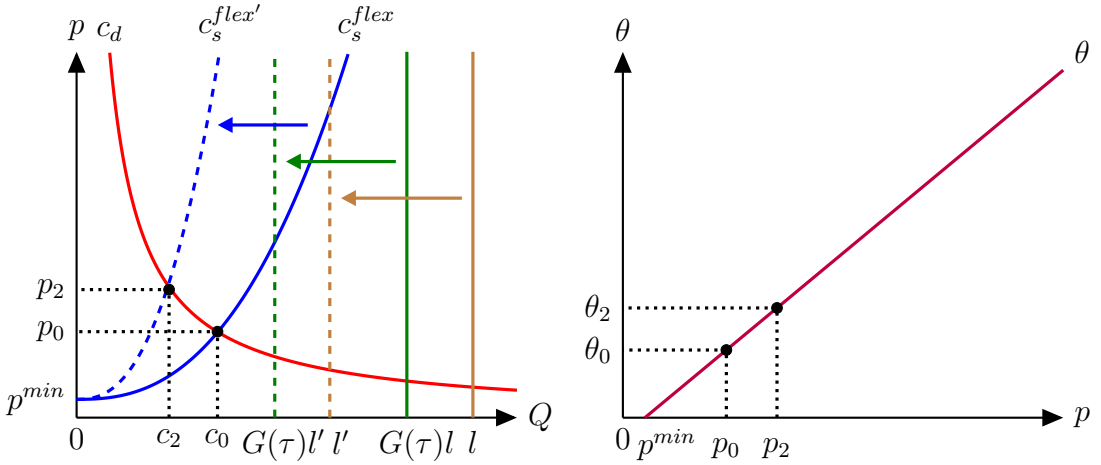

(b) An Adverse Shock to Productive Capacity, i.e., $l \downarrow$

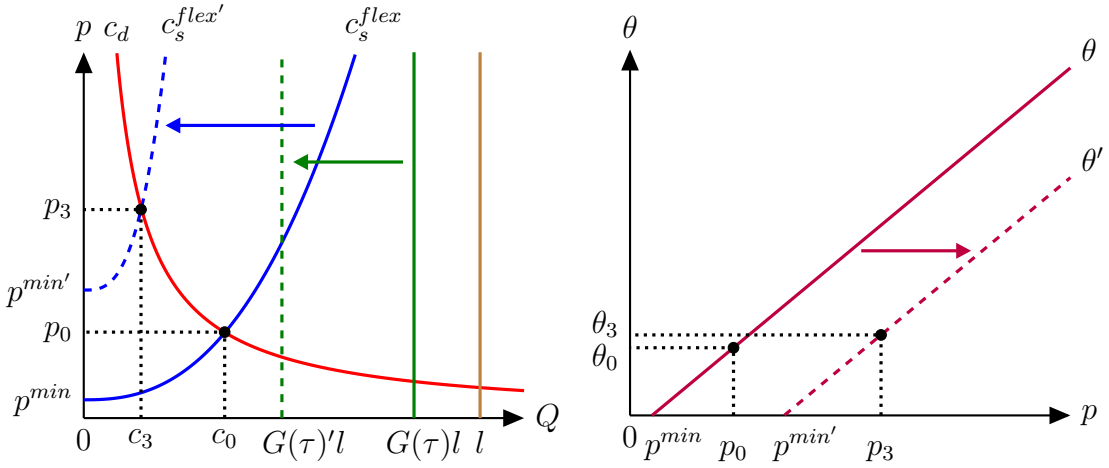

(c) An Adverse Shock to Supply Chain, i.e., $\gamma \uparrow$

Figure 9: Graphical Representation of Comparative Statics

and consumption falls. The higher price attracts more retailers to enter the market, which raises product market tightness. The simultaneous rise in price and product market tightness leads to a higher wholesale price. Matching costs and spare capacity also fall accordingly.

Panel (c) in Figure 9 shows the comparative statics for an increase in the scale parameter γ of the log-normal distribution of transportation costs $G(\cdot)$. A rise in the distribution of transportation costs increases the likelihood that producers will draw a transportation cost above the fixed reservation threshold. As a result, the fraction of goods shipped from producers decreases, curtailing the supply of goods available to the retailers and the households, and causing the price to increase to clear the market. Thus, price rises while consumption falls. Graphically, this process is represented by an inward shift of the aggregate supply curve from c_s^{flex} to $c_s^{flex'}$, together with an increase in p^{min}. Since the productive capacity of the economy remains unchanged while the number of successful trades falls, spare capacity increases.

The disturbance to the supply chain can either tighten or loosen product market tightness, depending on the extent to which the price rise compares to the fall in expected total surplus due to higher transportation costs. To consider these countervailing forces more explicitly, we revisit the equilibrium condition for product market tightness in Equation (27):

$$\theta(\gamma) = \frac{1-\eta}{\eta\rho}\left[\underbrace{p(\gamma)-\tau}_{\text{Profit Margin }\uparrow} + \underbrace{\beta\int_0^{\tau}\Phi\left(\frac{\log z'-\gamma}{\sigma}\right)dz'}_{\text{Expected Total Surplus }\downarrow}\right],$$

where $\Phi(\cdot)$ is the standard normal cumulative density function. This equation illustrates how the sensitivity of price to transportation costs determines the "profit margin," which incentivizes retailers to search for producers, thereby increasing tightness in the product market. Conversely, higher transportation costs decrease the "expected total surplus" from trading, which deters retailers from searching for producers. The net change in product market tightness resulting from a supply chain disruption is determined by the interplay of these two opposing forces.[28]

Important to our analysis in Section 5 on the effectiveness of monetary policy in controlling

[28]The simulation results in Appendix D.10 corroborate this point. Specifically, as illustrated in Table D.5, in the scenario of a 10% increase in γ, the resulting price increase is sufficient to raise product market tightness in the new steady state. However, in an alternative scenario with only a 0.0001% increase in γ, as shown in Table D.7, the resulting price increase does not lead to greater tightness in the product market.

inflation and output, changes in product market tightness also have implications for the sensitivity of goods supply to price variations following an adverse shock to the supply chain. Suppose the price increase is sufficiently large so that the rise in visits made by unmatched retailers significantly outpaces the increase in the number of unmatched producers (i.e., the goods market is congested on the retailers' side). While this imbalance leads to a tighter product market and increases the likelihood of a producer participating in trade, an additional retailer has only a limited impact on the producer's probability of forming a match. To see even a slight further increase in such a probability, prices would need to rise considerably more due to the diminishing returns to searching inherent in the constant-returns-to-scale matching function (see Equation (3)). Consequently, the number of matches and the supply of goods become less sensitive to price changes resulting from supply chain disruptions when the goods market is already tight. Graphically, this change in the sensitivity of goods supply to price variations is represented by a steeper slope of the aggregate supply curve, as depicted in Panel (c) of Figure 9.[29]

4. The Causal Effects of Supply Chain Disruptions

We are now ready to study the causal effects of supply chain disruptions by developing an SVAR model that utilizes our ACR index and constrains the responses of the macro aggregates to the three distinct shocks in line with our theoretical results from the model.[30]

Our empirical specification is based on Rubio-Ramírez et al. (2010) and Arias et al. (2018):

$$\boldsymbol{y}_t'\boldsymbol{A}_0 = \sum_{l=1}^{L} \boldsymbol{y}_{t-l}'\boldsymbol{A}_l + \boldsymbol{\omega}_t'\boldsymbol{C} + \boldsymbol{\varepsilon}_t', \quad 1 \leq t \leq T, \tag{28}$$

where $\boldsymbol{y}_t$ is an $n \times 1$ vector of the endogenous variables, $\boldsymbol{\omega}_t = [1, t]'$ is a 2×1 vector of a constant and a linear trend, $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of the structural shocks, $\boldsymbol{A}_l$ is an $n \times n$ matrix of structural

[29] Appendix D.4 demonstrates that the changes in the wholesale price and matching cost depend on the responses of product market tightness and price to the supply chain shock. Lastly, in Appendix D.5, we show that the slope of the aggregate supply curve is inversely related to product market tightness. Specifically, the aggregate supply curve becomes steeper in the (Q, p) plane as product market tightness increases. Figure D.1 illustrates the alternative scenario where the price increase is insufficient to raise product market tightness.

[30] In principle, we could also undertake a full structural estimation of our model. However, that strategy would require us to buy into too many ancillary assumptions (e.g., parametric forms, persistence of shocks, etc.), which can be problematic given the current state of knowledge about models of the global supply chain. While we used some of these assumptions to derive our identification restrictions, we are cautiously optimistic that the restrictions will hold for more general specifications, even if we can only demonstrate it numerically. Thus, this seems to be a situation where the additional flexibility offered by SVARs is most convenient.

parameters for $0 \leq l \leq L$ with $\boldsymbol{A}_0$ invertible, $\boldsymbol{C}$ is a $2 \times n$ matrix of parameters, L is the lag length, and T is the sample size. The vector $\boldsymbol{\varepsilon}_t$, conditional on past information and the initial conditions $\boldsymbol{y}_0, \ldots, \boldsymbol{y}_{1-L}$, is Gaussian with mean zero and variance-covariance matrix $\mathbf{1}_{n\times n}$, i.e., the $n \times n$ identity matrix. The SVAR model described in Equation (28) can be written compactly as follows:

$$\boldsymbol{y}_t' \boldsymbol{A}_0 = \boldsymbol{x}_t' \boldsymbol{A}_+ + \boldsymbol{\varepsilon}_t', \quad 1 \leq t \leq T, \tag{29}$$

where $\boldsymbol{A}_+' = [\boldsymbol{A}_1' \ \ldots \ \boldsymbol{A}_L' \ \boldsymbol{C}']$ and $\boldsymbol{x}_t' = [\boldsymbol{y}_{t-1}' \ \ldots \ \boldsymbol{y}_{t-L}' \ \boldsymbol{\omega}_t']$. The dimension of $\boldsymbol{A}_+$ is $m \times n$, where $m = nL + 2$. The reduced-form representation implied by Equation (29) is given by:

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

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

Motivated by the variables present in our theoretical model, we estimate our SVAR model using the monthly U.S. series for real GDP, personal consumption expenditures (PCE) goods price, and import price, alongside two constructed measures of spare capacity and product market tightness, as well as our ACR index, over the sample period from January 2017 to September 2023. All series have been seasonally adjusted.[31]

In addition to the standard series of macro aggregates typically considered in the literature for estimating the causal effects of global supply chain disruptions (Benigno et al., 2022; Finck and Tillmann, 2022; De Santis, 2024; Finck et al., 2024), we construct a measure of spare productive capacity that aligns with our theoretical model and include it in the SVAR estimation. Specifically, we calculate the average spare capacity rate of the top five exporting countries to the U.S. – namely, Mexico, Canada, China, Germany, and Japan, which collectively account for more than 53% of goods imported into the U.S. This rate is weighted by the U.S. imports of goods from each country, and defined as:

$$SpareCapacityRate_t = \sum_{i \in \mathscr{C}} \left[\frac{Import_{i,t}}{\sum_{i \in \mathscr{C}} Import_{i,t}} \cdot \left(1 - CapacityUtilization_{i,t}\right) \right], \tag{30}$$

where $\mathscr{C} \equiv \{Mexico, Canada, China, Germany, Japan\}$, $Import_{i,t}$ denotes the U.S. imports of goods by customs basis from country i in month t, and $CapacityUtilization_{i,t}$ represents the

[31] In Appendix E.7 we add the federal funds rate to the SVAR estimation and adjusting the identification restrictions accordingly. Thus, instead of identifying a general demand shock, we also consider a contractionary monetary policy shock. Since the main results are unchanged, we relate this more involved specification to the Appendix.

capacity utilization rate for country i in month t. Furthermore, in line with our theoretical model, we construct a measure of product market tightness by dividing the total U.S. manufacturers' new orders by the import-weighted total spare capacity of the aforementioned U.S. trading partners.[32] The U.S. dollar amount of our measure of spare capacity is estimated using the following formula:

$$SpareCapacityDollar_t = \sum_{i \in \mathscr{C}} \Bigg[\frac{Import_{i,t}}{\sum_{i \in \mathscr{C}} Import_{i,t}} \cdot \left(\frac{1}{CapacityUtilization_{i,t}} - 1 \right) \cdot IndustrialProduction_{i,t} \Bigg], \tag{31}$$

where $IndustrialProduction_{i,t}$ stands for the total industrial production of country i in month t, measured in U.S. dollars. Lastly, the import price is used as a proxy for the wholesale price in order to capture the international sourcing strategies of U.S. retailers, particularly during the COVID-19 pandemic.

All variables related to the U.S. are retrieved directly or constructed using available data from the Federal Reserve Economic Data (FRED), maintained by the Federal Reserve Bank of St. Louis.[33] For variables related to the top five exporting countries to the U.S., data are retrieved directly or computed using available data from the respective government authorities or statistics agencies.[34] Specifically, while the monthly time series for the capacity utilization rates for Mexico and Japan are readily available from official sources, those for Canada, China, and Germany are derived by interpolating the corresponding quarterly series using the Chow-Lin method for temporal disaggregation (Chow and Lin, 1971), based on the industrial production index of each country. Lastly, the monthly industrial production for each country, as referenced

[32] We use the total U.S. manufacturers' new orders to proxy for the number of unmatched retailers in the definition of product market tightness, a time series that is directly available from the FRED database. To ensure robustness, in Appendix E.4, we approximate U.S. retailers' new orders using their inventories and sales data. Using such an alternative series to calculate product market tightness yields quantitatively similar results to those using manufacturers' new orders.

[33] The mnemonics for the variables related to the U.S. are: `GDPC1` (real GDP), `INDPRO` (industrial production), `DGDSRG3M086SBEA` (PCE goods price), `IMPMX` (imports of goods from Mexico), `IMPCA` (imports of goods from Canada), `IMPCH` (imports of goods from China), `IMPGE` (imports of goods from Germany), `IMPJP` (imports of goods from Japan), `AMTMNO` (total manufacturers' new orders), and `IR` (import price). As in Bernanke and Mihov (1998) and Arias et al. (2019), the monthly time series for real GDP is constructed by interpolating the corresponding quarterly series based on the industrial production index. The monthly series for the PCE goods price and total manufacturers' new orders are raw series taken directly from FRED, while the series for imports of goods from each of the U.S.'s major trading partners and the import price are seasonally adjusted using the X-13ARIMA-SEATS algorithm.

[34] Specifically, data are obtained from the National Institute of Statistics and Geography for Mexico, Statistics Canada for Canada, the National Bureau of Statistics for China, the Ifo Institute for Economic Research for Germany, and the Ministry of Economy, Trade and Industry for Japan.

in Equation (31) and measured in U.S. dollars, is calculated using the corresponding monthly industrial production index and the total dollar value of goods produced in the manufacturing, mining, and utilities sectors in each country.

Real GDP, PCE goods price, product market tightness, and import price enter the SVAR model in log percent, while spare capacity and the ACR index enter in percent. We set the number of lags to two in the baseline specification, but the results are robust to considering longer lags.[35]

Our identification scheme applies the sign restrictions derived from our theoretical model, as summarized in Table 1, as well as the zero restrictions on the contemporaneous responses of the ACR index to adverse shocks to aggregate demand and productive capacity. We impose these zero restrictions to sharpen our identification of the structural shocks, and *more importantly*, because they are motivated by our domain knowledge of the shipping industry: as emphasized in Section 2.6, container ships will not alter their routes without at least several weeks' notice in response to shocks to aggregate demand or productive capacity, thereby leaving our ACR index largely independent of demand or capacity shocks within a month.[36] Nonetheless, we verify the robustness of our results by testing them without these restrictions in the estimation, as detailed in Appendix E.1. Also, we check in Appendix E.2 whether similar identification results can be achieved by omitting the ACR index from the estimation. The inclusion of both the ACR index and the zero restrictions imposed on its contemporaneous responses to demand and capacity shocks enables tighter identification results.

We estimate the SVAR model using the Bayesian approach as in Arias et al. (2018, 2019, 2023) with restrictions only on the first period of impulse responses (i.e., horizon $k = 1$), thus imposing a minimal structure as in Mumtaz and Zanetti (2012, 2015).[37] Furthermore, as illustrated in

[35] Appendix E.3 demonstrates the robustness of our baseline results when considering various lag structures, specifically one, three, or four lags. Additionally, we make several substitutions: real GDP is replaced with the real PCE of goods and the industrial production index; the PCE goods price with the GDP deflator; the import price index with the producer price index; and manufacturers' new orders with retailers' new orders for constructing an alternative measure of product market tightness. As shown in Appendix E.4, despite these substitutions, the results remain consistent.

[36] An alternative would be to estimate an SVAR model where the ACR index is interpreted as a noisy signal of the latent variable "state of the global supply chain." This could be done with a state space representation, where an SVAR that includes the "state of the global supply chain" as an unobserved variable serves as the transition equation, and the measurement equation links the "state of the global supply chain" with our ACR index. While this approach seems a worthwhile exercise, we leave it for future research, as it would require a more thorough treatment than what we can provide in this already lengthy study.

[37] We use a Normal-Generalized-Normal (NGN) prior distribution over $\boldsymbol{A}_0$ and $\boldsymbol{A}_+$. The NGN prior is a conjugate prior characterized by four parameters $(\nu, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\Omega})$. The parameters ν and $\boldsymbol{\Phi}$ govern the marginal

Appendix D.10, the convergence of the dynamic version of our theoretical model from one steady state to another following each shock of interest occurs almost instantaneously, and the process is monotonic. This supports our choice to impose identification restrictions only at horizon $k = 1$. More concretely, we impose the following restrictions:

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

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

Restriction 3. ***An adverse shock to the supply chain*** *leads to a negative response of real GDP, as well as to a positive response of PCE goods price, spare capacity, and the ACR at $k = 1$.*

Restrictions 1 and 2 reflect our discussion in Section 2 about the rigidity of shipping schedules and contracts. The ACR index does not respond within the month when an aggregate demand or supply shock occurs because port arrivals take more than a month to react. Liner shipping companies typically need three to six months to adjust their routes in response to these shocks, and ships require several additional weeks to complete their new routes. In comparison, Restriction 3 says that, after a global supply chain shock, the ACR can respond at impact (although we are not imposing that it must; the parameters estimates will determine whether this is the case).

We leave the responses of product market tightness and import prices unrestricted, as these responses depend on the specific type of supply chain disruption shock in question (i.e., higher transportation costs vs. lower matching efficiency).[38]

Figures 10, 11, and 12 show our baseline results for the responses of the endogenous variables to an adverse shock to aggregate demand, productive capacity, and the supply chain, respectively.

prior distribution of $vec(\boldsymbol{A}_0)$, while $\boldsymbol{\Psi}$ and $\boldsymbol{\Omega}$ govern the prior distribution of $vec(\boldsymbol{A}_+)$, conditional on $\boldsymbol{A}_0$. We pick $\nu = 0, \boldsymbol{\Phi} = \boldsymbol{0}_{n\times n}, \boldsymbol{\Psi} = \boldsymbol{0}_{m\times n}$, and $\boldsymbol{\Omega}^{-1} = \boldsymbol{0}_{m\times m}$. This parameterization generates prior densities that are equivalent to those in Uhlig (2005). Appendix E.5 also ascertains that our results are robust to using the prior robust approach in Giacomini and Kitagawa (2021).

[38] For robustness, in an alternative SVAR estimation, we apply additional positive sign restrictions on the contemporaneous responses of product market tightness and import price to an adverse shock to the supply chain. This robustness check follows our theoretical prediction in Proposition 5 for a supply chain disturbance represented by a reduction in matching efficiency. As shown in Appendix E.6, the results are quantitatively similar to those obtained without imposing such restrictions.

The solid lines show the point-wise posterior median impulse response functions (IRFs) of the endogenous variables to each structural shock, and the gray-shaded areas represent the corresponding 68% and 90% posterior probability bands.[39] The shape and size of our estimated IRFs to demand and capacity shocks are comparable to those found in classic papers that estimate the responses of macro aggregates to traditional demand and supply shocks (Christiano et al., 1999; Peersman, 2005; Smets and Wouters, 2007; Fry and Pagan, 2011).

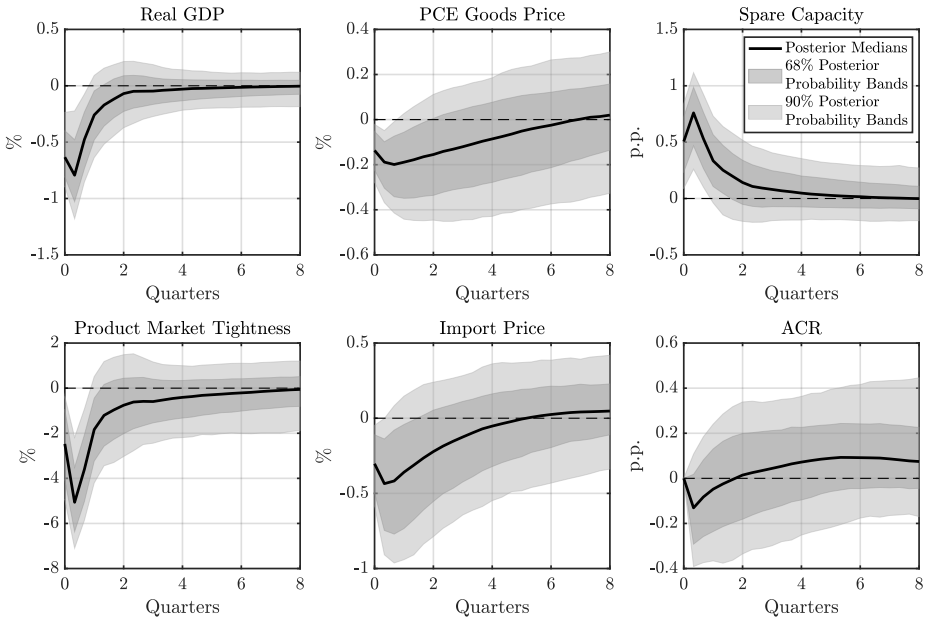

Figure 10: IRFs to an Adverse Shock to Aggregate Demand

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using the ACR index and Restrictions 1, 2, and 3. The solid line shows the point-wise posterior medians and the shaded bands represent the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

We begin by discussing the IRFs to an adverse shock to aggregate demand, as shown in Figure 10. On impact, real GDP declines significantly by more than 0.5%, and spare capacity rises sharply by approximately 0.5%. These responses persist with a high posterior probability for the first six months following the shock. Product market tightness initially falls substantially by about 2.5%, continues to drop by an additional 2.5%, and then gradually reverts to zero. In contrast, the response of the PCE goods price is muted, with an initial decrease of about 0.2% before gradually reverting to zero after eighteen months. The import price exhibits a similar pattern but returns to baseline more quickly, at the five-quarter mark. Lastly, the response of the

[39] These results are based on 100,000 independent draws from the posterior distribution of the structural parameters, with the structural shocks normalized to one standard deviation.

ACR index is negative shortly after impact but is less precisely estimated, with a large posterior probability mass centered around zero.

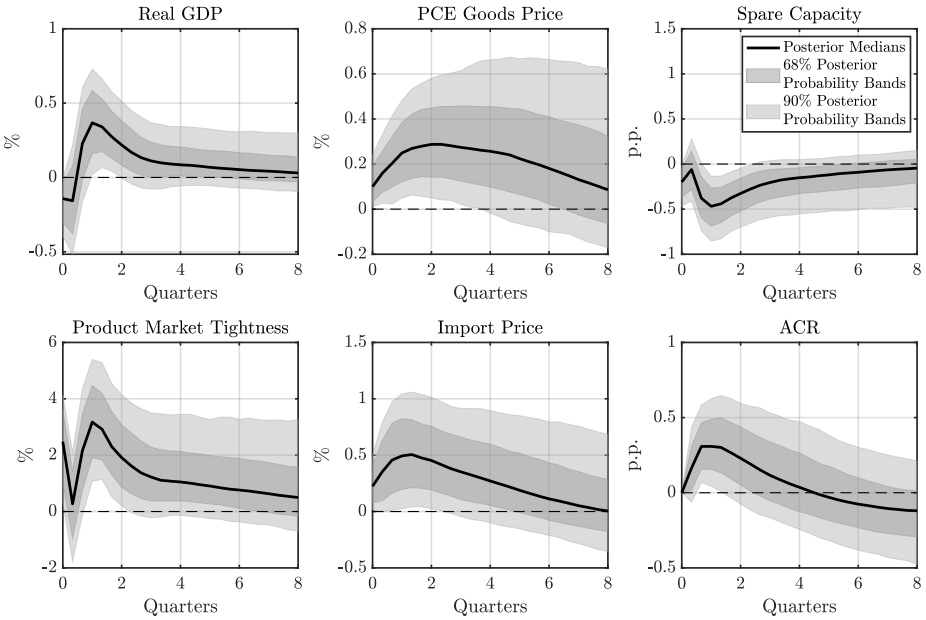

Figure 11: IRFs to an Adverse Shock to Productive Capacity

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using the ACR index and Restrictions 1, 2, and 3. The solid line shows the point-wise posterior medians and the shaded bands represent the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

Figure 11 shows the IRFs to an adverse shock to productive capacity. On impact, real GDP and spare capacity respond negatively, while the responses of the PCE goods price, import price, and product market tightness are positive, in accordance with Restriction 2. Immediately post-impact, despite a significant initial increase in product market tightness, the initial drop in productive capacity moderates the impacts on real GDP (recall the definition of aggregate supply in Equation (19)) and on spare capacity. Moreover, the rise in product market tightness is also short-sighted, as the PCE goods and import prices have not yet risen sufficiently to attract more unmatched retailers into the market, thus highlighting the lagged effects of supply-side disruptions. Subsequently, as inflationary pressures continue to build, the influx of retailers increases, leading to a marked increase in product market tightness by the one-quarter mark. Consequently, the initial decline in real GDP quickly reverses, with a rise of approximately 0.5% within one quarter of the shock. Spare capacity further declines by about 0.3% after the initial drop. Lastly, the response of the ACR index is constrained to zero on impact but peaks at approximately 0.3%

after the one-quarter mark, before gradually returning to zero. This pattern underscores the potential complementarities in disruptions to both the production and transportation processes of goods.

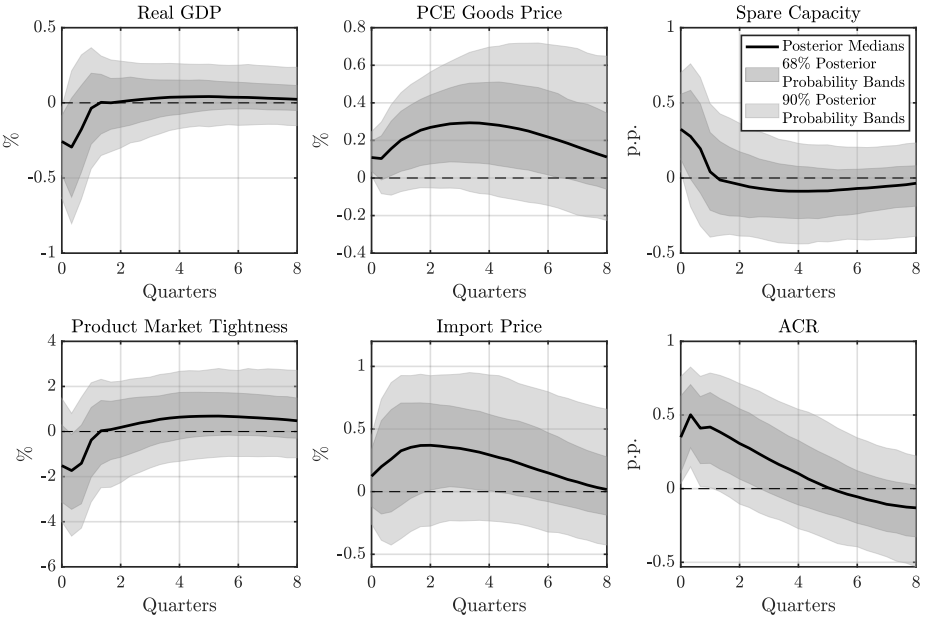

Figure 12: IRFs to an Adverse Shock to Supply Chain

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using the ACR index and Restrictions 1, 2, and 3. The solid line shows the point-wise posterior medians and the shaded bands represent the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

Figure 12 shows the IRFs following an adverse shock to the supply chain. The median response of real GDP is negative on impact and stays below zero for over a quarter after the shock. While real GDP decreases, spare capacity increases by roughly the same magnitude and remains elevated for one quarter. In terms of the response of product market tightness, which is unrestricted, it initially decreases before gradually reverting to zero.[40] The surges in both the PCE goods and import prices are consistent with the magnitudes observed following the negative capacity shock, highlighting the substantial impact of supply chain disruptions on price inflation.[41] Despite the uncertainty around our estimates, as indicated by the wide posterior probability bands, the

[40] In line with our theoretical prediction in Section 3.4, the initial decrease in tightness can be largely attributed to the rise in spare capacity following the supply chain disruption when prices have not adjusted. Subsequently, as prices continue to rise, more retailers are drawn into the product market, resulting in increased tightness.

[41] We also test whether similar identification results for the supply chain disturbance can be obtained using local projections with controls (Barnichon and Brownlees, 2019). This approach is equivalent to a recursive identification scheme where the ACR index is placed first in a recursive ordering, implying that all other endogenous variables can respond contemporaneously to shocks to the ACR index. As illustrated in Appendix E.8, the responses of the endogenous variables exhibit patterns similar to those shown in Figure 12.

positive responses of both the PCE goods and import prices remain within the 68% probability band. Furthermore, for the PCE goods price, the lower boundary of the 68% probability band approaches the zero-response line as the corresponding median response peaks at approximately the one-year mark. In contrast, the peak response for the import price arrives earlier, just two quarters post-impact. Lastly, the ACR index remains elevated for five quarters after the shock.

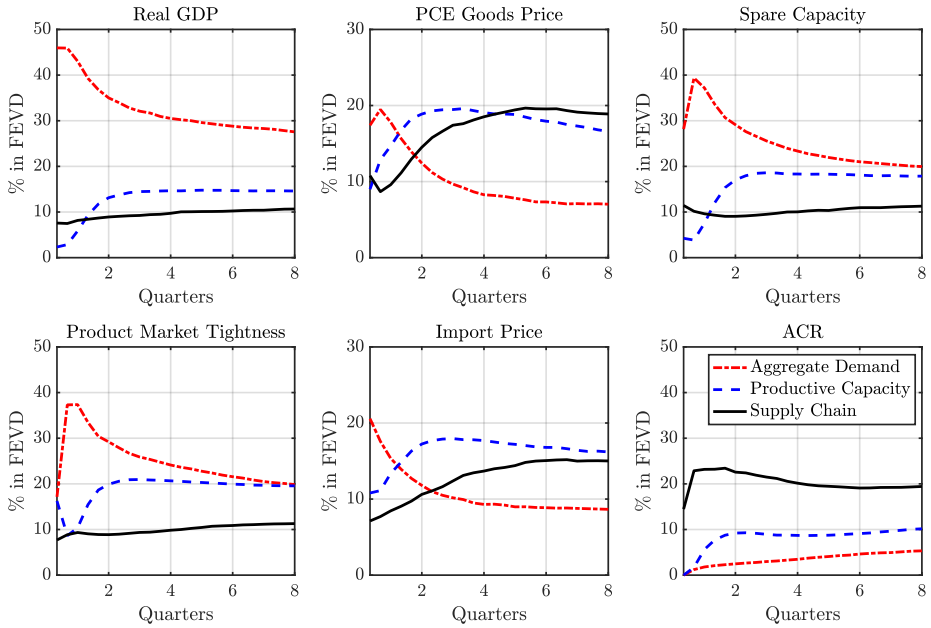

Figure 13: Forecast Error Variance Decomposition (FEVD) from the SVAR

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated using the ACR index and Restrictions 1, 2, and 3, and based on 100,000 independent importance sampling draws.

Figure 13 shows the proportion of forecast error variance explained by each of the three structural shocks. The aggregate demand shock accounts for the largest share of unexpected fluctuations in real GDP, spare capacity, and product market tightness across all horizons. However, its dominance diminishes over the long term as capacity shocks become more significant. Conversely, while the demand shock explains the majority of unexpected variations in the PCE goods and import prices at shorter horizons, capacity and supply chain disturbances grow more influential in the medium and long terms. They account for a larger portion of the unexpected fluctuations in both price indicators over these extended periods, indicating that supply-side disruptions have lasting effects on price dynamics. Furthermore, capacity shocks are more significant in explaining the unexpected variations in real GDP, spare capacity, and product market tightness than supply chain shocks. Lastly, supply chain shocks are the primary factor explaining the largest

proportion of unexpected fluctuations in the ACR index across all horizons. In contrast, demand and capacity shocks play only a negligible role in the short term, consistent with observations that port congestion is minimally affected by changes in demand and productive capacity over short-term horizons, as discussed in Section 2.

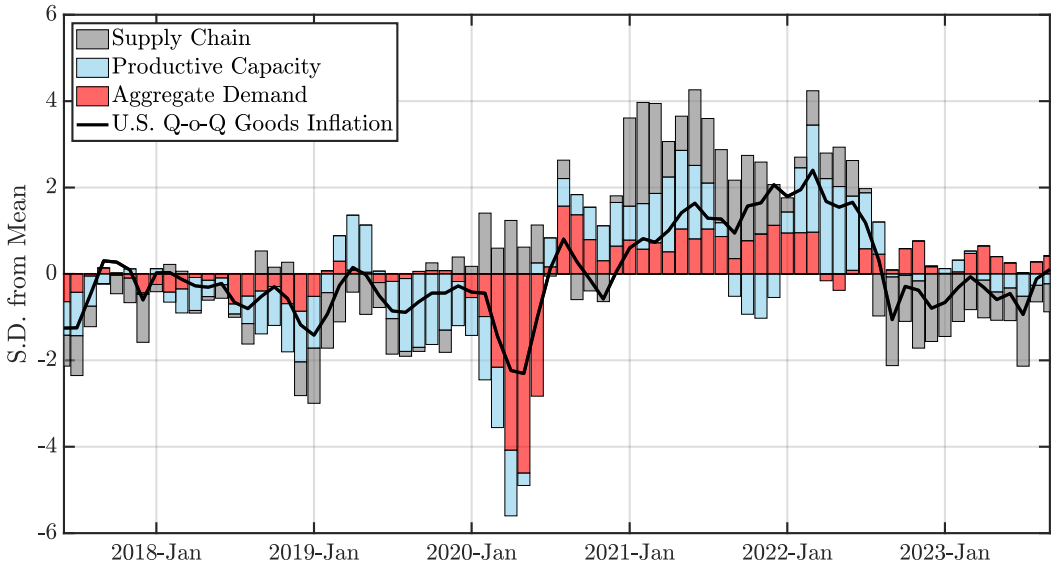

Figure 14: Historical Decomposition (HD) of U.S. Quarter-on-Quarter Goods Inflation

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with the ACR index included as the measure indicative of the state of the global supply chain, and with Restrictions 1, 2, and 3 imposed on the contemporaneous IRFs. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

Figure 14 displays the key empirical finding of our analysis. It shows the cumulative historical contribution of each of the three structural shocks to U.S. quarter-on-quarter goods inflation for the sample period from June 2017 to September 2023.[42] Using our ACR index and the theory-predicted identification restrictions, our estimation can be summarized in five findings.

First, supply chain disturbances consistently generated a negative contribution to inflation

[42]To facilitate the comparison of series across different scales, we have applied Z-score standardization, which rescales data to have a mean of zero and a standard deviation of one. Moreover, the accumulation of historical contributions begins in June 2017, as the estimation of these contributions starts in April 2017 due to the inclusion of two lags in the estimation, and the fact that the endogenous variables enter the SVAR model in first differences to facilitate the calculation of the growth rates. Lastly, to improve clarity, we leave out historical contributions from the lag structure, a constant, and other unidentified structural shocks in Figure 14. Consequently, the historical contributions from the three identified structural shocks, depicted by the shaded bars, do not exactly sum to the U.S. quarter-on-quarter goods inflation rate, which is represented by the black solid line.

prior to the start of the COVID-19 pandemic.[43] Second, the initial drop in inflation at the onset of the COVID-19 pandemic in early 2020 was largely attributed to a substantial decrease in aggregate demand, likely linked with mobility restrictions (which lowered the desire to consume) and elevated uncertainty. Third, the subsequent rises in inflation, especially those during 2021, were mainly due to adverse shocks to the supply chain. Fourth, the landscape changed in the first half of 2022: adverse shocks to productive capacity kept inflation elevated. Evidence from the major trading partners of the U.S. during this period supports our findings. For instance, in China, economic disruptions caused by outbreaks of the Omicron variant and the associated strict COVID-related mobility restrictions across major cities and provinces severely impacted the labor market and discouraged investment, thereby significantly reducing China's production possibilities (World Bank, 2022). Furthermore, higher unemployment benefits post-pandemic (e.g., the continued use of the Kurzarbeit scheme in Germany) and a shift in worker preferences toward more flexible work arrangements have played a pivotal role in reshaping the labor markets in the economies that trade extensively with the U.S., putting pressure on the productive capacities of these economies and, at the same time, keeping U.S. inflation elevated in the first half of 2022. Interestingly, aggregate demand played a relatively small role in driving inflation, suggesting that monetary and fiscal policy might not have been excessively expansionary. Fifth, from the second half of 2022 onward, a combination of weakened demand, strengthened capacity, and supply chain recovery has driven down inflation.

Table 2 also reports the cumulative historical contribution of each shock to U.S. annual goods inflation, highlighting that aggregate demand shocks were the main driving force behind the dynamics of inflation in 2020, supply chain shocks in 2021, productive capacity shocks in 2022, and the convolution of all three shocks in driving inflation from 2023 onward.[44]

[43] This finding supports the notion of strategic enhancements to supply chain operations to alleviate inflationary pressures. For instance, several U.S. ports (e.g., Port of Los Angeles) underwent considerable infrastructure upgrades between 2017 and 2019, aiming to increase their capacity, efficiency, and resilience against potential systematic disruptions.

[44] In Appendix F, we present the estimation results using alternative indices in the existing literature that track the state of the global supply chain, including the HARPEX, the New York Fed's GSCPI, and the SDI by Smirnyagin and Tsyvinski (2022) and Liu et al. (2024), among others. The main takeaway is that our ACR index yields significantly tighter identification results compared to these alternatives in terms of the estimated causal effects of supply chain disturbances.

Table 2: Cumulative Historical Contribution of Each Shock to U.S. Goods Inflation

Date	U.S. Goods Inflation (Percent)	Cumulative Historical Contribution		
		Aggregate Demand (Percent)	Productive Capacity (Percent)	Supply Chain (Percent)
Apr-Dec 2017	-0.04	-0.33	-0.39	-0.30
Jan-Dec 2018	-0.02	-0.59	-1.08	-0.07
Jan-Dec 2019	0.37	-0.04	-0.78	-0.22
Jan-Dec 2020	-0.12	-1.01	0.13	0.18
Jan-Dec 2021	8.16	0.96	0.67	1.11
Jan-Dec 2022	4.67	0.40	1.56	-0.32
Jan-Sep 2023	1.21	0.36	-0.19	-0.36

Notes. Each row shows the U.S. goods inflation rate, calculated as the growth of the PCE goods price index, along with the cumulative historical contributions of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation for each sample year from 2017 to 2023. Note that 2017 and 2023 are not fully represented years due to a lack of data. The shocks are identified using the SVAR specification in Equation (28), with the ACR index included as the measure indicative of the state of the global supply chain, and with Restrictions 1, 2, and 3 imposed on the contemporaneous IRFs. The numbers reported for the cumulative historical contributions are the posterior medians, based on 100,000 independent importance sampling draws.

5. The Effectiveness of Monetary Policy

Our next task is to study the interplay between supply chain disruptions and the effectiveness of monetary policy in controlling inflation and output. First, we show through our theoretical model that a disruption to the supply chain increases the sensitivity of inflation and reduces the sensitivity of output to a contractionary monetary policy shock, generating state-dependence in the stabilization trade-off for monetary policy. Then, we will test and empirically corroborate our theoretical prediction using a threshold vector autoregression (TVAR) model.

5.1. Theoretical Prediction

We derive the theoretical prediction for the state-dependence of monetary policy by returning to our model in Section 3. The money supply parameter μ encapsulates the action of monetary policy, and the scale parameter of the log-normal distribution of transportation costs γ captures the disruption to the supply chain (in the interest of space, the case where the matching efficiency falls is relegated to Appendix D.9; suffice it to say that we get the same results). We study the comparative statics of the impacts of a tightening in monetary policy, focusing on whether the

effects of the policy intervention on inflation and output are different amid the supply chain disruption.[45] Proposition 6 summarizes our results.

Proposition 6. *For any given threshold of reservation transportation cost $\tau > 0$ and parameter values relevant for monetary policy $\mu \in \mathbb{R}^+$ and transportation costs $\gamma \in \mathbb{R}$, when product market tightness is sufficiently elevated to allow producers to recoup the increase in transportation costs due to the supply chain disruption, as represented by the following constraint:*

$$\frac{\partial \theta(\mu,\gamma)}{\partial \gamma} > \frac{\theta(1+\theta^{\xi})}{(1-G(\tau))\,G(\tau)} \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(\log\tau-\gamma)^2}{2\sigma^2}\right], \tag{32}$$

where $G(\tau) \equiv \Phi[(\log\tau-\gamma)/\sigma]$, $\Phi(.)$ is the standard normal cumulative density function, the responses of the endogenous variables to a change in monetary policy are described by the partial derivatives:

$$\frac{\partial c(\mu,\gamma)}{\partial \mu} > 0,\ \frac{\partial p(\mu,\gamma)}{\partial \mu} > 0,\ \frac{\partial \theta(\mu,\gamma)}{\partial \mu} > 0,\ \frac{\partial r(\mu,\gamma)}{\partial \mu} > 0,$$
$$\frac{\partial}{\partial \mu}\left[G(\tau)l - c(\mu,\gamma)\right] < 0,\ \frac{\partial}{\partial \mu}\left[l - c(\mu,\gamma)\right] < 0.$$

The cross derivatives of the endogenous variables that describe the optimal interplay between a change in monetary policy and the supply chain disruption are given by:

$$\frac{\partial^2 c(\mu,\gamma)}{\partial \mu \partial \gamma} < 0,\ \frac{\partial^2 p(\mu,\gamma)}{\partial \mu \partial \gamma} > 0,\ \frac{\partial^2 \theta(\mu,\gamma)}{\partial \mu \partial \gamma} > 0,\ \frac{\partial^2 r(\mu,\gamma)}{\partial \mu \partial \gamma} > 0,$$
$$\frac{\partial^2}{\partial \mu \partial \gamma}\left[G(\tau)l - c(\mu,\gamma)\right] > 0,\ \frac{\partial^2}{\partial \mu \partial \gamma}\left[l - c(\mu,\gamma)\right] > 0,$$

where $c, p, \theta, r, G(\tau)l-c$, and $l-c$ represent consumption (or, equivalently, output), price, product market tightness, wholesale price, matching cost, and spare capacity, respectively.

Proof. See Appendix D.6. ■

The partial and cross derivatives in Proposition 6 imply that when the increase in product market tightness is sufficiently large during the supply chain disruption, and therefore producers have greater incentives to trade with retailers (as stated in Equation (32)), the supply chain

[45] In Appendix D.7, we derive the theoretical prediction for the effectiveness of monetary policy, depending on whether the productive capacity of the economy is constrained or not. Similar to the scenario in which the supply chain is disrupted, contractionary monetary policy is more effective at taming inflation and reducing the sensitivity of output when the productive capacity is constrained. The only difference is that the state-dependent effects of monetary policy are unconditional.

disruption intensifies the fall in inflation while dampening the fall in consumption (or, equivalently, output) that is associated with a contractionary monetary policy shock.[46]

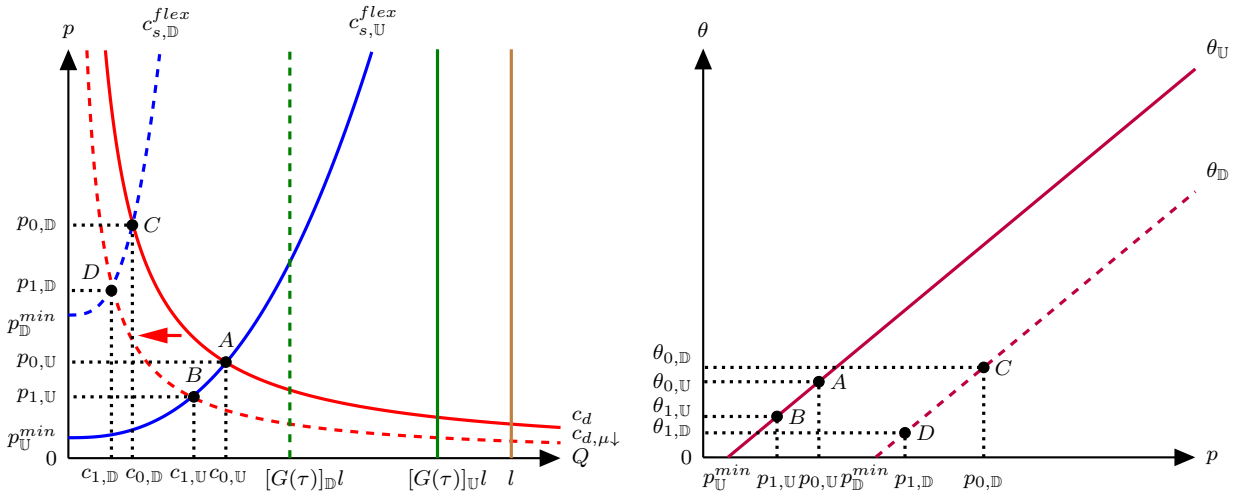

Figure 15: State-Dependent Effects of a Contractionary Monetary Policy Shock: Theoretical Prediction

Notes. The panels illustrate the adjustment of the economy to a contractionary monetary policy shock and to what extent the adjustment depends on the supply chain disruption. The two respective states – i.e., supply chain disrupted ($\mathbb{D}$) vs. undisrupted ($\mathbb{U}$) – are plotted against each other. $c_{s,\mathbb{D}}^{flex}$ and $c_{s,\mathbb{U}}^{flex}$ represent the aggregate supply curves in the two states, while $\theta_{\mathbb{D}}$ and $\theta_{\mathbb{U}}$ represent the schedules of equilibrium product market tightness in the two states. c_d and $c_{d,\mu\downarrow}$ denote the aggregate demand curves before and after the contractionary monetary policy shock, respectively. The labels on the axes corresponding to each state are differentiated by their subscripts, and the capital letters ($A \to B, C \to D$) indicate the dynamics of the economy in the two states.

Figure 15 provides the graphical representation of our theoretical prediction. In response to a contractionary monetary policy shock, households reduce consumption due to decreased money holdings. This reduction causes the aggregate demand curve to shift inward, leading to lower prices and reduced consumption of goods. Consequently, product market tightness decreases as retailers visit fewer producers to purchase goods. This reduction in demand and product market tightness leads to a lower wholesale price. Facing diminished demand, producers sell a smaller fraction of their productive capacity to retailers, which results in decreased output. Consequently, matching costs and spare capacity increase.

Recall from our discussion in Section 3.4 that the aggregate supply curve becomes steeper when the increase in product market tightness during the supply chain disruption is sufficiently large, as described in Equation (32). This steepening occurs because the probability of producers engaging in trade becomes less sensitive to price changes when the market is already tight. In

[46] Proposition 6 also shows that during the supply chain disruption, the responses of product market tightness and wholesale price are more pronounced, while the responses of matching cost and spare capacity are less pronounced.

such scenarios, the number of matches is constrained by the shorter side, namely, the number of unmatched producers. As a result, the supply of goods becomes less responsive to price changes during the supply chain disruption. Consequently, a contractionary monetary policy shock significantly reduces inflation with only a relatively modest decrease in output.

5.2. Empirical Validation

We test our theoretical prediction for the state-dependence of monetary policy by developing a structural TVAR model – building on Chen and Lee (1995) – that allows for endogenous variations in the parameters based on the estimated threshold of our ACR index. The reduced-form TVAR model is:

$$\boldsymbol{y}_t = I_t \left[\sum_{l=1}^{L} \boldsymbol{B}'_{\mathbb{D},l} \boldsymbol{y}_{t-l} + \boldsymbol{C}'_{\mathbb{D}} \boldsymbol{\omega}_t + \boldsymbol{\Sigma}_{\mathbb{D}}^{1/2} \boldsymbol{\varepsilon}_t \right] + (1 - I_t) \left[\sum_{l=1}^{L} \boldsymbol{B}'_{\mathbb{U},l} \boldsymbol{y}_{t-l} + \boldsymbol{C}'_{\mathbb{U}} \boldsymbol{\omega}_t + \boldsymbol{\Sigma}_{\mathbb{U}}^{1/2} \boldsymbol{\varepsilon}_t \right], \tag{33}$$

where $1 \leq t \leq T$, $\boldsymbol{y}_t$ is an $n \times 1$ vector of the endogenous variables, $\boldsymbol{\omega}_t = [1, t]'$ is a 2×1 vector of a constant and a linear trend, $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of the structural shocks, $\boldsymbol{B}_{\mathbb{D},l}$ and $\boldsymbol{B}_{\mathbb{U},l}$ are two $n \times n$ matrices of coefficients for the lagged endogenous variables $\boldsymbol{y}_{t-l}$, $\boldsymbol{C}_{\mathbb{D}}$ and $\boldsymbol{C}_{\mathbb{U}}$ are two $2 \times n$ matrices of coefficients for the constant and linear trend, $\boldsymbol{\Sigma}_{\mathbb{D}}$ and $\boldsymbol{\Sigma}_{\mathbb{U}}$ are the variance-covariance matrices (we allow the variance-covariance matrix to be regime-specific), L is the lag length, and T is the sample size. The vector $\boldsymbol{\varepsilon}_t$, conditional on past information and the initial conditions $\boldsymbol{y}_0, \ldots, \boldsymbol{y}_{1-L}$, is Gaussian with mean zero and variance-covariance matrix $\mathbf{1}_{n \times n}$, i.e., an $n \times n$ identity matrix. Switches between the two regimes – i.e., supply chain disrupted ($\mathbb{D}$) vs. undisrupted ($\mathbb{U}$) – are governed by the indicator variable $I_t \in \{0, 1\}$, which is equal to one if the ACR index in period $t-1$, ACR_{t-1}, is above the threshold $\overline{ACR}$, and equal to zero otherwise:[47]

$$I_t = \begin{cases} 1, & \text{if } ACR_{t-1} > \overline{ACR}; \\ 0, & \text{if } ACR_{t-1} \leq \overline{ACR}. \end{cases} \tag{34}$$

Under the Normal-Inverse-Wishart conjugate prior for the TVAR parameters and conditional on the value of the threshold $\overline{ACR}$, the posterior distribution of the TVAR parameter vector is a conditional Normal-Inverse-Wishart distribution, and we use the Gibbs sampler to draw from the distribution. We use a Metropolis-Hastings algorithm to obtain the posterior distribution

[47] In principle, we could have different threshold levels, but two regimes are sufficient to document the argument we are making.

of the threshold $\overline{ACR}$ conditional on the TVAR parameters, similar to Chen and Lee (1995), Lopes and Salazar (2006), and Pizzinelli et al. (2020). Appendix G.1 provides the details on the Normal-Inverse-Wishart prior.

To retain comparability with our previous empirical results, we include the same variables used in our SVAR model in Section 4, with the addition of the federal funds rate to reflect changes in the stance of U.S. monetary policy. For consistency, we also retain the same sample period from January 2017 to September 2023.[48]

To identify the contractionary monetary policy shock, we follow the theoretical prediction presented in Proposition 6 and impose the following standard restriction on the IRFs:

Restriction 4. ***A contractionary monetary policy shock*** *leads to a negative response of real GDP, PCE goods price, product market tightness, and import price, as well as to a positive response of spare capacity and the federal funds rate at* $k = 1$*. The ACR does not respond at* $k = 1$*.*[49]

We compute the identified set of IRFs using the Bayesian approach similar to that in Pizzinelli et al. (2020) and Bratsiotis and Theodoridis (2022).[50] We use one lag in the baseline estimation and select the one-month lag of the ACR index as the variable that determines the state I_t.[51]

Figure 16 plots the IRFs to a contractionary monetary policy shock for both the supply chain disrupted (black) and undisrupted (red) regimes, reporting both the point-wise posterior medians (solid lines) and the 68% equal-tailed point-wise posterior probability bands (shaded area and dotted lines) from horizon $k = 0$ up to horizon $k = 12$ (i.e., four quarters). The figure shows significant differences in the responses of the endogenous variables to the contractionary monetary policy shock between the two regimes. In accordance with Proposition 6, the PCE goods and import prices are more responsive, while real GDP and spare capacity are less responsive in the

[48] All series have been seasonally adjusted, except the federal funds rate. Real GDP, PCE goods price, product market tightness, and import price enter the TVAR in log percent, whereas the federal funds rate, spare capacity, and the ACR index enter in percent.

[49] Restriction 4 enriches Restriction 1, which is intended for the identification of an adverse shock to aggregate demand, by including the positive response of the federal funds rate on impact, which is the main instrument to control monetary policy.

[50] To implement the sign and zero restrictions on the IRFs, we use the penalty function approach (PFA) developed in Uhlig (2005) and Mountford and Uhlig (2009). The PFA consists of using a loss function to find an orthogonal matrix that satisfies the zero restrictions and that satisfies or comes close to satisfying the sign restrictions. Appendix G.2 provides the details on the PFA.

[51] Appendix G.3 plots the posterior distribution of the threshold $\overline{ACR}$, together with the time series of the identified regimes using the median of the posterior $\overline{ACR}$.

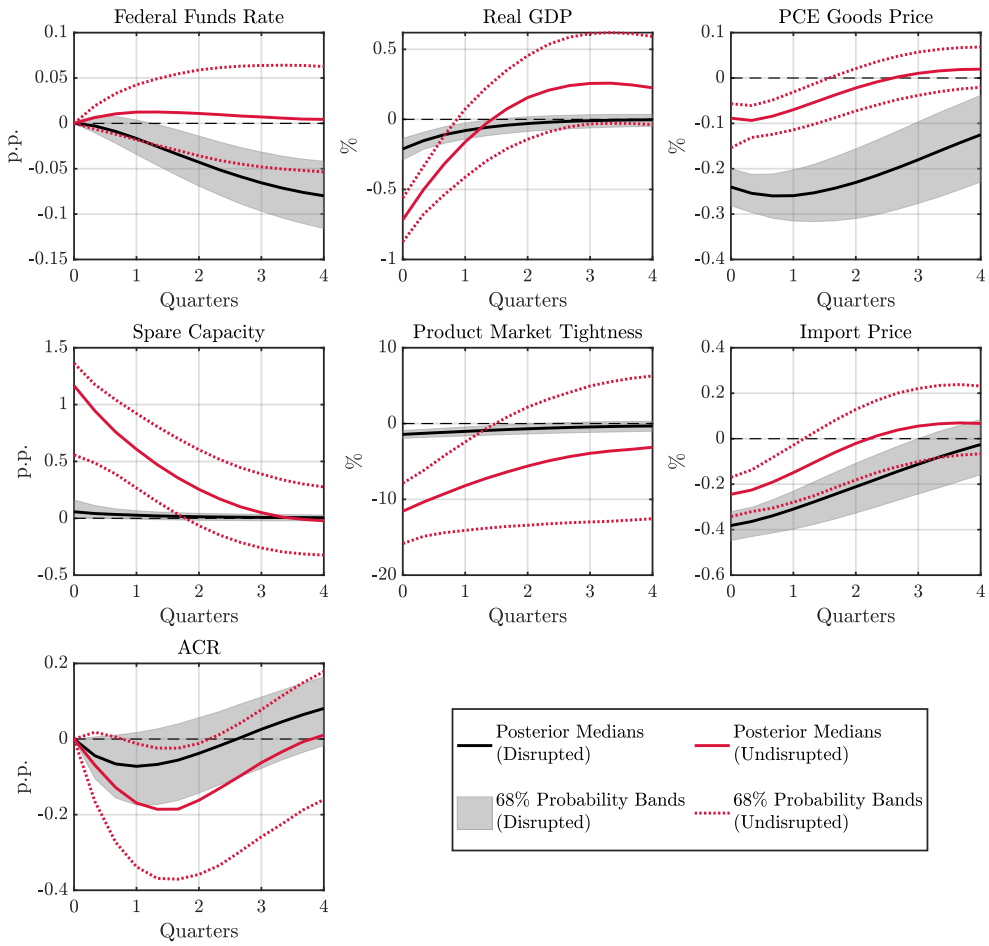

Figure 16: State-Dependent Effects of a Contractionary Monetary Policy Shock: Empirical Validation

Notes. The figure shows the IRFs to a one standard deviation contractionary monetary policy shock identified using Restriction 4 for both the supply chain disrupted and undisrupted regimes. The solid black (solid red) line shows the point-wise posterior medians, and the shaded black area (dotted red lines) depicts the 68% equal-tailed point-wise posterior probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 10,000 independent draws from the posterior.

regime where the supply chain is disrupted, and the differences in the responses are statistically significant. The responses of the federal funds rate in the two regimes also align with the observed patterns for prices and output, as the federal funds rate remains elevated throughout the horizons in the undisrupted regime while dipping below zero in the disrupted regime. The responses of product market tightness between the two regimes, however, are contrary to our theoretical prediction, with a much more muted response observed for the disrupted regime, suggesting other mechanisms at play in shaping the imbalance between supply and demand in the product market. Lastly, the responses of the ACR index between the two regimes cannot be disentangled from each other.

Appendix H shows that our results hold across several variations to the benchmark model: (i) using the Wu-Xia shadow federal funds rate (Wu and Xia, 2016) to reflect the stance of U.S. monetary policy; (ii) dropping the zero restriction imposed on the contemporaneous response of the ACR index; (iii) employing different lag structures; and (iv) adopting a looser prior (i.e., $\lambda = 0.5$ instead of 0.25; see Appendix G.1 for details on the tightness of the prior). Appendix I also shows that our results continue to hold when we use local projections with interaction terms, as in Ramey and Zubairy (2018), Ghassibe and Zanetti (2022), and Arias et al. (2023), to estimate the state-dependent effects of a contractionary monetary policy shock.

6. Conclusion

Our study constructs the first index of global supply chain disruptions by analyzing congestion at major container ports worldwide using high-frequency satellite data from AIS, available publicly since 2017. We develop a novel machine learning-based spatial clustering algorithm that uses the locations, speeds, and headings of container ships to accurately distinguish between berth and anchorage areas within ports of various geographical morphologies, thereby allowing us to measure port congestion with high precision. By aggregating congestion rates across the top 50 container ports, we create the ACR index of global port congestion, providing a measure of the state of the global supply chain.

We develop a new theoretical framework that includes separate production and retailing processes, search frictions in the exchanges between producers and retailers, and transportation costs. The model simultaneously generates spare capacity for producers and a scarcity of supply for retailers, leading to sharp price increases and heightened search frictions that curtail output in response to disruptions to the supply chain. Our framework demonstrates that disturbances to the supply chain reduce output and increase prices, as in standard models, *and distinctively*, they also increase the spare productive capacity. The co-movements of spare productive capacity, prices, and output allow us to uniquely identify supply chain disruption shocks and study their causal effects on macroeconomic outcomes through an SVAR model with sign and zero restrictions. The empirical model establishes that supply chain disturbances lead to prolonged, substantial increases in prices and a detrimental effect on real GDP along with the simultaneous rise in the spare productive capacity.

We also show, both theoretically and empirically, that monetary policy exerts a stronger influence on inflation, albeit with a diminished effect on output, amid supply chain disruptions. Thus, disruptions to the supply chain enhance the effectiveness of contractionary monetary policy in taming inflation while simultaneously reducing the sensitivity of output to the policy.

Our study opens several important avenues for future research. First, our new index reveals significant heterogeneity in the congestion of ports around the world. It would be interesting to study whether the spillovers between ports are primarily driven by geographical proximity or the production synergies that Fernández-Villaverde et al. (2021, 2024a,b) find critical for trading relationships. The presence of heterogeneity raises the possibility of reducing port congestion arising from supply chain disturbances by strategically reorganizing the locations of trading partners in accordance with their production synergies. Second, our results show that spare productive capacity is central to the ramifications of supply chain disturbances in the economy. Thus, enriching the analysis by endogenizing the adjustment of spare capacity and studying its persistence might be important.[52] Third, it would be worthwhile to consider input-output networks and explore the role of spare capacity in the transmission of supply chain disruptions across firms in the network. The structure of the production network could potentially magnify or dampen the disturbances to the supply chain, which may trigger endogenous changes in the structure of the network, as documented in Ghassibe (2023). Fourth, the incorporation of predictive analytics into our spatial clustering algorithm will enable the algorithm to anticipate supply chain disruptions by identifying changes and systematic patterns in the shipments using real-time positions, speeds, and weights of container ships. The enriched algorithm could prove a powerful tool to design preemptive policy actions to offset or minimize the disruptions to the supply chain. We plan to pursue some of these extensions in our future work.

[52]Seminal studies by Ramey (1989), Burnside et al. (1995), and Basu (1996) suggest an important role for inventories and spare capacity in business cycle fluctuations.

References

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

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

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

Antràs, P. (2023). An 'Austrian' Model of Global Value Chains. *NBER Working Paper 30901*.

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

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

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

Attinasi, M. G., Balatti, M., Mancini, M., and Metelli, L. (2021). Supply Chain Disruptions and the Effects on the Global Economy. *ECB Economic Bulletin Issue 8*.

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

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

Baqaee, D. and Farhi, E. (2022). Supply and Demand in Disaggregated Keynesian Economies With an Application to the COVID-19 Crisis. *American Economic Review*, 112:1397–1436.

Barnichon, R. and Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *The Review of Economics and Statistics*, 101(3):522–530.

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

Basu, S. (1996). Procyclical Productivity: Increasing Returns or Cyclical Utilization? *Quarterly Journal of Economics*, 111(3):719–751.

Bekaert, G., Engstrom, E., and Ermolov, A. (2020). Aggregate Demand and Aggregate Supply Effects of COVID-19: A Real-Time Analysis. *Finance and Economic Discussion Series 2020-049*.

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

Benigno, P. and Eggertsson, G. B. (2023). It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. *NBER Working Paper 31197*.

Benigno, P. and Ricci, L. A. (2011). The Inflation-Output Trade-off With Downward Wage Rigidities. *American Economic Review*, 101(4):1436–66.

Bernanke, B. S. and Mihov, I. (1998). Measuring Monetary Policy. *Quarterly Journal of Economics*, 113:869–902.

Blanchard, O. J. and Bernanke, B. S. (2023). What Caused the US Pandemic-Era Inflation? *NBER Working Paper 31417*.

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

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

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

Bratsiotis, G. J. and Theodoridis, K. (2022). Precautionary Liquidity Shocks, Excess Reserves and Business Cycles. *Journal of International Financial Markets, Institutions and Money*, 77:101518.

Brown, J., Englert, D., and Hoffmann, J. (2021). International Transport Costs: Why and How to Measure Them? World Bank Blogs – Transport. Available at: https://blogs.worldbank.org/transport/international-transport-costs-why-and-how-measure-them (Accessed: March 6, 2022).

Bureau of Transportation Statistics (2021). On National Maritime Day and Every Day, U.S. Economy Relies on Waterborne Shipping. Data Spotlights. Available at: https://www.bts.gov/data-spotlight/national-maritime-day-and-every-day-us-economy-relies-waterborne-shipping (Accessed: May 15, 2024).

Burnside, C., Eichenbaum, M., and Rebelo, S. (1995). Capital Utilization and Returns to Scale. *NBER Macroeconomics Annual*, 10:67–110.

Chen, C. W. and Lee, J. C. (1995). Bayesian Inference of Threshold Autoregressive Models. *Journal of Time Series Analysis*, 16:483–492.

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

Chow, G. C. and Lin, A.-l. (1971). Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series. *Review of Economics and Statistics*, 53:372–375.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (1999). Chapter 2 Monetary Policy Shocks: What Have We Learned and to What End? In *Handbook of Macroeconomics*, volume 1, pages 65–148. Elsevier.

Comín, D. A., Johnson, R. C., and Jones, C. J. (2023). Supply Chain Constraints and Inflation. *NBER Working Paper 31179*.

Costinot, A., Vogel, J., and Wang, S. (2013). An Elementary Theory of Global Supply Chains. *Review of Economic Studies*, 80:109–144.

Coşar, A. K. and Demir, B. (2018). Shipping Inside the Box: Containerization and Trade. *Journal of International Economics*, 114:331–345.

De Santis, R. A. (2024). Supply Chain Disruption and Energy Supply Shocks: Impact on Euro-Area Output and Prices. *International Journal of Central Banking*, 20(2):193–235.

den Haan, W. J., Ramey, G., and Haefke, C. (2005). Turbulence and Unemployment in a Job Matching Model. *Journal of the European Economic Association*, 3:1360–1385.

di Giovanni, J., Şebnem Kalemli-Özcan, Silva, A., and Yildirim, M. A. (2022). Global Supply Chain Pressures, International Trade, and Inflation. *NBER Working Paper 30240*.

di Giovanni, J., Şebnem Kalemli-Özcan, Silva, A., and Yildirim, M. A. (2023). Pandemic-Era Inflation Drivers and Global Spillovers. *NBER Working Paper 31887*.

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

Eichenbaum, M., Rebelo, S., and Wong, A. (2022). State-Dependent Effects of Monetary Policy: The Refinancing Channel. *American Economic Review*, 112(3):721–61.

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

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

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

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

Fernández-Villaverde, J., Yu, Y., and Zanetti, F. (2024b). Technological Synergies, Heterogeneous Firms, and Idiosyncratic Volatility. *NBER Working Paper 32247*.

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

Finck, D. and Tillmann, P. (2022). The Macroeconomic Effects of Global Supply Chain Disruptions. *BOFIT Discussion Paper 14/2022*.

Freightos (2024). Shipping From Shanghai, Shanghai to Los Angeles, CA: Air, Sea & Container Freight. Freightos. Available at: https://www.freightos.com/routes/route/cn-shanghai-shanghai/us-losangeles-ca (Accessed: May 22, 2024).

Fry, R. and Pagan, A. (2011). Sign Restrictions in Structural Vector Autoregressions: A Critical Review. *Journal of Economic Literature*, 49(4):938–60.

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

Fujita, S. (2018). Declining Labor Turnover and Turbulence. *Journal of Monetary Economics*, 99:1–19.

Ganapati, S., Wong, W. F., and Ziv, O. (2021). Entrepôt: Hubs, Scale, and Trade Costs. *NBER Working Paper 29015*.

Ghassibe, M. (2023). Endogenous Production Networks and Non-Linear Monetary Transmission. *CREi Working Paper Series*.

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

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

Gordon, M. V. and Clark, T. E. (2023). The Impacts of Supply Chain Disruptions on Inflation. *Federal Reserve Bank of Cleveland Economic Commentary 2023-08*.

Hall, R. E. (2005). Employment Fluctuations With Equilibrium Wage Stickiness. *American Economic Review*, 95(1):50–65.

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

Heiland, I., Moxnes, A., Ulltveit-Moe, K. H., and Zi, Y. (2022). Trade From Space: Shipping Networks and the Global Implications of Local Shocks. *CEPR Discussion Paper 14193*.

Howitt, P. and McAfee, R. P. (1987). Costly Search and Recruiting. *International Economic Review*, 28(1):89–107.

Ikeda, D., Li, S., Mavroeidis, S., and Zanetti, F. (2024). Testing the Effectiveness of Unconventional Monetary Policy in Japan and the United States. *American Economic Journal: Macroeconomics*, 16(2):250–286.

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

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

Keynes, J. M. (1940). *How to Pay for the War*. Macmillan and Co.

Kharroubi, E. and Smets, F. (2024). Monetary Policy with Profit-Driven Inflation. *CEPR Discussion Paper 18946*.

Kim, S.-J. and Shin, H. S. (2023). Theory of Supply Chains: A Working Capital Approach. *BIS Working Paper 1070*.

Lane, P. R. (2022). Bottlenecks and Monetary Policy. ECB Blog. Available at: https://www.ecb.europa.eu/press/blog/date/2022/html/ecb.blog220210~1590dd90d6.en.html (Accessed: February 10, 2022).

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

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

Liu, P., Theodoridis, K., Mumtaz, H., and Zanetti, F. (2019). Changing Macroeconomic Dynamics at the Zero Lower Bound. *Journal of Business & Economic Statistics*, 37(3):391–404.

Lopes, H. F. and Salazar, E. (2006). Bayesian Model Uncertainty in Smooth Transition Autoregressions. *Journal of Time Series Analysis*, 27:99–117.

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

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

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

Morison, S. E. (1954). *The Atlantic Battle Won: May 1943-May 1945*. Little, Brown and Company.

Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *Review of Economic Studies*, 61(3):397–415.

Mountford, A. and Uhlig, H. (2009). What Are the Effects of Fiscal Policy Shocks? *Journal of Applied Econometrics*, 24:960–992.

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

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

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

Notteboom, T., Pallis, T., and Rodrigue, J. (2021). Disruptions and Resilience in Global Container Shipping and Ports: The COVID-19 Pandemic Versus the 2008–2009 Financial Crisis. *Maritime Economics & Logistics*, 23:179–210.

OECD and EUIPO (2021). *Misuse of Containerized Maritime Shipping in the Global Trade of Counterfeits*. OECD.

Peersman, G. (2005). What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions. *Journal of Applied Econometrics*, 20(2):185–207.

Pizzinelli, C., Theodoridis, K., and Zanetti, F. (2020). State Dependence in Labor Market Fluctuations. *International Economic Review*, 61:1027–1072.

Ramey, V. A. (1989). Inventories as Factors of Production and Economic Fluctuations. *American Economic Review*, 79(3):338–354.

Ramey, V. A. and Zubairy, S. (2018). Government Spending Multipliers in Good Times and in Bad: Evidence From US Historical Data. *Journal of Political Economy*, 126:850–901.

Ramondo, N. and Rodríguez-Clare, A. (2013). Trade, Multinational Production, and the Gains From Openness. *Journal of Political Economy*, 121:273–322.

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

Samuelson, P. A. (1954). The Transfer Problem and Transport Costs, II: Analysis of Effects of Trade Impediments. *The Economic Journal*, 64(254):264–289.

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

Smirnyagin, V. and Tsyvinski, A. (2022). Macroeconomic and Asset Pricing Effects of Supply Chain Disasters. *NBER Working Paper 30503*.

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

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

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

The White House (2021). Fact Sheet: Biden Administration Efforts to Address Bottlenecks at Ports of Los Angeles and Long Beach: Moving Goods from Ship to Shelf. The White House. Available at: https://www.whitehouse.gov/briefing-room/statements-releases/2021/10/13/fact-sheet-biden-administration-efforts-to-address-bottlenecks-at-ports-of-los-angeles-and-long-beach-moving-goods-from-ship-to-shelf/ (Accessed: October 13, 2021).

The World Bank (2022). *Global Economic Prospects, June 2022*. Washington, DC: World Bank.

Thomas, C. and Zanetti, F. (2009). Labor Market Reform and Price Stability: An Application to the Euro Area. *Journal of Monetary Economics*, 56(6):885–899.

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

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

UNCTAD (2019). Review of Maritime Transport 2019. UNCTAD. Available at: https://unctad.org/webflyer/review-maritime-transport-2019 (Accessed: November 1, 2021).

World Bank (2022). China Economic Update – June 2022: Between Shocks and Stimulus – Real Estate Vulnerabilities and Financial Stability in China. Technical report, World Bank Group, Washington, DC.

Wu, J. C. and Xia, F. D. (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking*, 48(2-3):253–291.

Xu, L., Yu, Y., and Zanetti, F. (2024). The Adoption and Termination of Suppliers Over the Business Cycle. Technical report, University of Oxford, Department of Economics.

Online Appendices

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

Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li, Francesco Zanetti

Contents

A Background on the Containerized Shipping Industry **A-3**

B A Density-Based Spatial Clustering Algorithm **A-5**

B.1 Methodology A-5

B.2 Illustrative Case: Port of Ningbo-Zhoushan A-15

B.3 Weekly Indices of Port Congestion A-22

C Discussion on the Assumptions in the Model **A-24**

D Long Proofs and Model Dynamics **A-27**

D.1 Proof of Proposition 1 . . . . . A-27

D.2 Proof of Proposition 2 . . . . . A-28

D.3 Proof of Proposition 4 . . . . . A-29

D.4 Proof of Proposition 5 . . . . . A-29

D.5 Slope of the Aggregate Supply Curve and Its Dependence on Product Market Tightness A-33

D.6 Proof of Proposition 6 . . . . . A-34

D.7 Theoretical Prediction on the Effectiveness of Monetary Policy When Productive Capacity Is Constrained A-37

D.8 Fixed Price Aggregate Supply A-40

D.9 Supply Chain Disruptions and Reduced Matching Efficiency A-43

D.10 Convergence Dynamics A-49

E Robustness of SVAR Results **A-58**

E.1 Dropping the Zero Restrictions A-58

E.2 Dropping the ACR Index A-62
E.3 Different Lag Structures A-66
E.4 Alternative Proxies for Consumption, Prices, and Product Market Tightness . . . A-71
E.5 Prior Robustness A-82
E.6 Additional Sign Restrictions on Product Market Tightness and Import Price . . . A-86
E.7 Adding the Federal Funds Rate A-90
E.8 Identification Using Local Projections A-94

F Alternative Indices Measuring the State of the Global Supply Chain A-97
F.1 HARPEX A-97
F.2 GSCPI A-101
F.3 SDI A-106
F.4 Other Indices in the Literature A-109
F.5 ACR for the Major Ports Along the Trans-Pacific Route A-110
F.6 ACT A-113

G Priors and Identification in the TVAR A-116
G.1 Priors A-116
G.2 Identification Using the PFA A-118
G.3 Posterior and Identified Regimes A-119

H Robustness of TVAR Results A-121

I State-Dependence Results Using Local Projections A-126

References for Appendices A-134

A. Background on the Containerized Shipping Industry

In the appendix, we provide additional background on the containerized shipping industry, highlighting several important features related to its short-term rigidity in response to prevailing economic conditions.

Port congestion and speed adjustment. As mentioned in the main text, the shipping industry generally adopts a "hurry up and wait" practice regarding the port call process (Du et al., 2015). For instance, a vessel departs the loading port at full speed, aiming to meet the originally requested time of arrival at the pilot boarding place (RTA PBP) scheduled for day 14. However, if the port encounters severe congestion three days into the voyage, altering the RTA PBP to day 17, the ship may not receive this updated information in time to adjust its speed. Even if the ship does receive the forewarnings, it often chooses not to alter its speed, as such adjustments might violate contractual obligations. Consequently, even when a port is experiencing congestion, vessels will still "hurry" to arrive and then "wait" at anchorage. This "hurry up and wait" phenomenon is common in the container shipping industry, where operational inefficiencies arise from mismatches between shipping schedules and real-time port conditions.

Oil price, speed adjustment, and congestion. Fuel costs account for approximately 50% to 60% of a vessel's operating costs for a liner shipping company (Notteboom, 2006). Moreover, the fuel consumption of a vessel is roughly a cubic function of the sailing speed (Li et al., 2016). Hence, vessel sailing speed significantly impacts operating costs. In principle, shipping companies would dynamically adjust sailing speeds based on current bunker oil prices. However, based on AIS data, researchers have found that this relationship between speed and oil prices is not apparent in practice (Adland and Jia, 2016). Furthermore, by regressing the monthly average speed of container ships on the West Texas Intermediate (WTI) crude oil futures price after taking the natural logarithm of both series, we find that the two series are statistically uncorrelated, with an estimated coefficient of -0.0035 and a robust standard error of 0.0078, yielding a p-value of 0.657. The corresponding R^2 stands at 0.0031. This has been attributed to the rigidity of contractual structures and the lack of coordination between ports and vessels or, simply, that ships are optimized for operating at one particular speed (with changes in oil prices affecting the design of the next generation of ships, e.g., fuel-efficient cargo vessels). Consequently, the observed impact of oil prices on vessel speed, and thus on port congestion, is limited.

Idle ships. According to the Clarksons Shipping Intelligence Network, idle container ships are defined as vessels not recorded with an average speed greater than one knot for seven days or more, not identified as subject to another status (e.g., laid-up, under repair, in storage, or similar), and not subsequently recorded with an average speed greater than one knot for two or more consecutive days or not having moved more than 20 km.

Figure A.1 plots the series of idle ships expressed as a percentage of the entire global fleet. This proportion declined from 6% to approximately 3% during 2017, then hovered around 4% through 2018–2019 before soaring to slightly below 9% at the onset of the COVID-19 pandemic in early 2020. Subsequently, it dived back down and stabilized at 4% thereafter. As illustrated in Section 2.6, the spike in the proportion of idle ships at the onset of the pandemic is closely related to the active capacity management by shipping companies, which set aside capacity in response to the unprecedented declines in consumer and business demand. To minimize the effect of such an abrupt change in the occurrence of idling on the estimation of port congestion, we exclude idle ships in the construction of the congestion indices (see Appendix B.1.1).

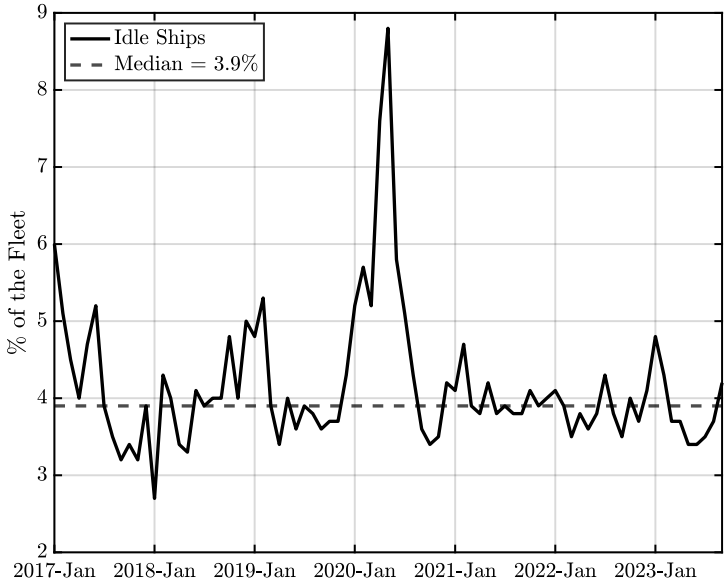

Figure A.1: Idle Ships

Notes. Proportion of container ships with an idle status in the global fleet. Idle status applies to ships not recorded with an average speed greater than one knot for seven days or more, not identified as subject to another status (e.g., laid-up, under repair, in storage, or similar), and either not subsequently recorded with an average speed greater than one knot for two or more consecutive days or not having moved more than 20 km.

B. A Density-Based Spatial Clustering Algorithm

In this appendix, we provide the technical details of our density-based spatial clustering algorithm, namely the iterative, multi-attribute, density-based spatial clustering of applications with noise (IMA-DBSCAN). Most of the technical details provided in this appendix can also be found in the companion paper, Bai et al. (2023).

This algorithm is used to estimate port congestion for the top 50 container ports worldwide.[1] In subsequent sections, we first delve into the methodology underpinning our algorithm. We then present an illustrative case where we apply the algorithm to the Port of Ningbo-Zhoushan in China, demonstrating its capability to identify both anchorage and berth areas of a port, where other methods fall short. Lastly, we present the weekly congestion indices – namely, the ACR and the ACT – to highlight that our estimation of port congestion is robust to using different time frequencies.

B.1. Methodology

As depicted in Figure B.1, the proposed IMA-DBSCAN algorithm has several distinct features. Foremost among these is its two-tiered iterative structure. At the first level, we extract the trajectory of each container ship at each of our 50 ports from the AIS data. For each ship, a traditional DBSCAN (Ester et al., 1996) is employed to filter out noise and cluster all its mooring points. While this level can pinpoint mooring areas, it does not adequately differentiate between anchorage and berth areas of a port. The second level addresses this limitation. Here, we apply a spatial-temporal DBSCAN (ST-DBSCAN; see Birant and Kut, 2007) to the clustering. During this phase, we employ an iterative method to determine a generalized and optimal parameter setting for the clustering algorithm. Another hallmark of IMA-DBSCAN is its integration of multiple attributes at the second level. Beyond spatial data, such as coordinates, we also weave in non-spatial information, such as headings and timestamps, to enhance clustering accuracy. Next, we elaborate on the specifics of each level of IMA-DBSCAN.

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

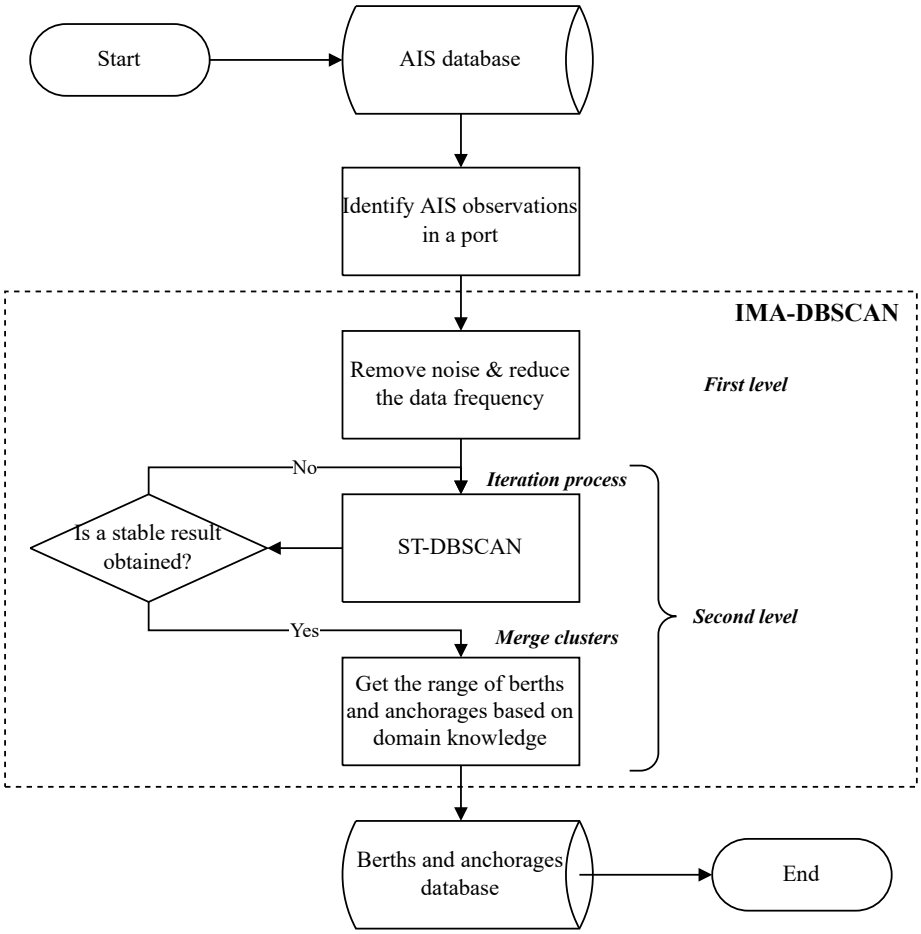

Figure B.1: Methodology Framework of IMA-DBSCAN

B.1.1. The First Level – Data Preprocessing

While AIS data provide detailed information on the positions of each ship, directly clustering these positions to determine the anchorage and berth areas of a port presents several challenges. First, even if we restrict the data to a specific port area within a certain timeframe, the sheer volume of records means that inputting them directly into DBSCAN would result in extended processing times. Second, a high incidence of incorrect AIS signal assignments could lead to inaccurate clustering outcomes, such as identifying a cluster that is not an actual berth or one that covers an unusually large geographical area. Third, if a ship stays in a port area for an extended period, the dense AIS data could cause DBSCAN to mistakenly identify it as a cluster. Given these challenges, we must preprocess the AIS data.

In the first level of IMA-DBSCAN, we begin by filtering the AIS data for each ship in the port area, focusing on records indicating speeds of less than one knot. Such positions suggest that a

ship is either berthed, anchored, or in an unusual situation (e.g., idle, laid-up, under repair, in storage, or similar). We then tally these positions; if their number falls outside an acceptable range (e.g., less than 100 or more than 100,000), we deem the ship's data to be abnormal and exclude it from further analysis. Since a ship might dock at a port multiple times, we establish a period, Δt (e.g., 12 hours), as the cutoff between two consecutive arrivals. If the gap between two arrivals exceeds Δt, we treat them as separate port calls. To streamline the data while maintaining consistency, we retain only the first data point for each hour. For every port call of a ship, its positions are clustered using the traditional DBSCAN with parameters Eps and $MinPts$. We choose an Eps value small enough to identify the ship's mooring areas and an appropriate $MinPts$ value to ensure transient stops are classified as noise. At this stage, the preprocessing of AIS data is complete. The refined samples are then used to identify the anchorage and berth areas of a port in the second level of IMA-DBSCAN. For reference, the pseudo-code for the first level of IMA-DBSCAN is detailed in Algorithm 1.

B.1.2. The Second Level – Multiple Attributes and Iteration

Information on headings. As highlighted in the main text, AIS data integrate both spatial (i.e., geographical coordinates) and non-spatial (i.e., headings) information. In Figure 1 in the main text, we illustrate the positions of a ship in a port alongside its headings. We observe that the headings of a ship at a berth are either aligned in the same direction or are exact opposites. In contrast, headings in an anchorage area appear random, with no discernible pattern. This observation aligns with real-world scenarios, where ships in anchorage areas often struggle to maintain consistent headings over time due to significant wind and wave variations.

Consequently, in the second level of IMA-DBSCAN, we leverage this heading information to enhance estimation accuracy.[2] Specifically, IMA-DBSCAN incorporates three parameters: $Eps1$, $Eps2$, and $MinPts$. Here, $Eps1$ denotes the maximum geographical coordinate (spatial) distance, $Eps2$ represents the maximum non-spatial distance between two headings, and $MinPts$ is the minimum number of points within the distances defined by $Eps1$ and $Eps2$. The geographical

[2]Such non-spatial information is also useful when distinguishing between different berths (see Algorithm 2). In our initial experiment, the coordinates could only help us identify the approximate locations of anchorage and berth areas of a port, not the exact number of berths.

coordinate (spatial) distance, D, is calculated using the Haversine formula:

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

where the geographical coordinates are measured in radians and $R = 6,371$ is the mean radius of Earth in kilometers. On the other hand, the non-spatial distance Δh between two headings is calculated as follows:

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

With the two measures of distance defined above, the neighbors of a point are those with a geographical coordinate (spatial) distance less than $Eps1$ and a non-spatial distance less than $Eps2$. A core is defined as a point with a number of neighbors greater than or equal to $MinPts$. The clusters in IMA-DBSCAN consist only of these core points.

Iteration process. Given that the geographical shapes of anchorage and berth areas vary significantly across ports, and the boundaries of anchorage areas of each port are constantly changing due to unpredictable waves and winds, the values of the three parameters in IMA-DBSCAN should ideally vary to achieve optimal estimation results. Hence, we propose an iterative method to determine these parameter values. Specifically, while we fix $Eps2$ at 1°, our method allows the values of $Eps1$ and $MinPts$ to vary between different ports. During the iteration process, we define four intermediate variables: $Dist$, m, m', and $NumC$. Here, $Dist$ represents the average distance between a point in a cluster and the center of that cluster, m denotes the number of points, and m' represents the number of noisy points, which is initialized to zero. Lastly, $NumC$ indicates the number of clusters.[3] Using these intermediate variables, $MinPts$, and $Eps1$ are calculated as follows:

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

Regarding α and β, even though there is no explicit constraint on their values, they should fall within a reasonable range to ensure both the convergence of our algorithm and the validity

[3]Since there are no clusters at initialization, we treat all points as if they were part of the same cluster. Additionally, if all points are classified as noise, we set $NumC = 1$.

of the identification results. After evaluating the performance of IMA-DBSCAN under various parameter settings, we find that an admissible range of $0.4 \leq \alpha \leq 0.6$ and $0.06 \leq \beta \leq 0.1$ is appropriate. We also introduce an intermediate variable, $Dist_0$, which records the value of $Dist$ from the previous iteration and is initialized to zero.

Following this, we execute ST-DBSCAN iteratively. In each iteration, ST-DBSCAN operates with $Eps1$ and $MinPts$ set to their current values, and $Eps2$ set to 1°. The outputs classify each point either into a cluster or as noise. Based on these outputs, the values of the intermediate variables, as well as those for $Eps1$ and $MinPts$, are updated. These updated values are then reapplied in ST-DBSCAN for the subsequent iteration. The entire process concludes when the difference, $Dist - Dist_0$, is less than or equal to $\Delta Dist$ (e.g., 100 m).[4] Consequently, each point is either assigned to a cluster or labeled as noise. We then interpret the areas of clusters as berths and the areas with noisy points as anchorages.

Information on timestamps. After running ST-DBSCAN, we find that a large proportion of clusters should be merged, as they essentially represent the same berth in reality. To achieve a more accurate identification of berth areas, we merge certain clusters by utilizing the timestamp information in the AIS data. More precisely, we first calculate the start and end times of each port call in each cluster. Subsequently, since only one ship can dock at a berth at any given moment, for each cluster under consideration, we identify the closest cluster and check for any overlap in docking times. If there is at least one overlap, the two clusters are considered to represent two different berths. If there is no overlap, the two clusters are merged to represent one berth (see Figure B.2 for an illustration).

Furthermore, to differentiate between anchorage areas of a port, we perform another DBSCAN on those points previously classified as noise. In this process, the two parameters associated with DBSCAN, i.e., Eps' and $MinPts'$, are set according to our domain knowledge. Finally, we remove clusters with fewer than N port calls, where N is also determined based on our domain knowledge. For reference, the pseudo-codes for the second level of IMA-DBSCAN can be found in Algorithms 2, 3, and 4.

Lastly, in estimating port congestion for the top 50 container ports worldwide, the parameter values set for IMA-DBSCAN are provided in Table B.1.

[4]Introducing a two-tiered structure and incorporating non-spatial information in clustering do not undermine the efficiency of IMA-DBSCAN. Specifically, the values of $Eps1$ and $MinPts$ stabilize after only five iterations, with additional iterations having minimal impact on the clustering results.

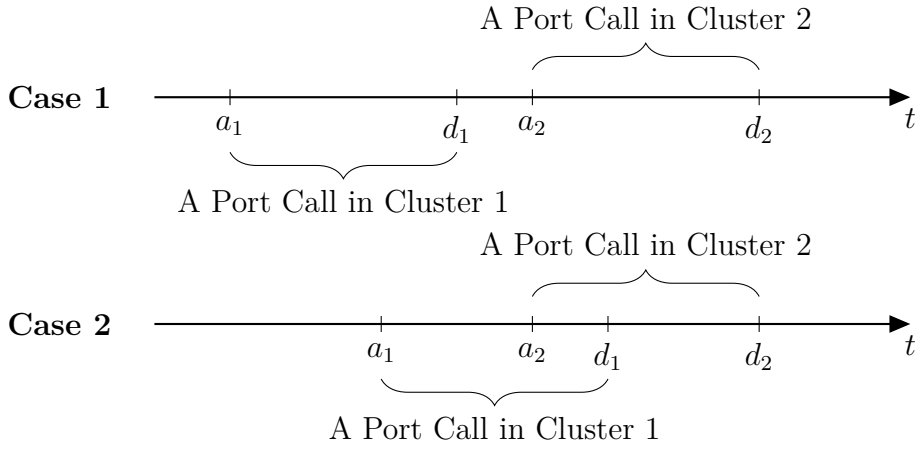

Figure B.2: Merging Clusters

Notes. The figure illustrates two scenarios and discusses the criteria for merging clusters after executing ST-DBSCAN at the second level. Here, a_1 and d_1 represent the arrival and departure times of a ship during a port call assigned to cluster 1. Similarly, a_2 and d_2 correspond to the times for a port call assigned to cluster 2, which is geographically the closest to cluster 1. In the first scenario, there is no overlap in the docking times, so clusters 1 and 2 are merged. In contrast, the second scenario shows an overlap in the docking times. As a result, clusters 1 and 2 are kept separate since two ships cannot occupy a single berth simultaneously.

Table B.1: Parameter Values for IMA-DBSCAN

	Parameter					
First Level	Δt	Eps	$MinPts$			
Value	12 hours	50 m	10			
Second Level	α	β	$\Delta Dist$	Eps'	$MinPts'$	N
Value	0.5	0.08	100 m	1,000 m	50	5

Algorithm 1 Level 1 IMA-DBSCAN

Inputs:

$A_l = \{a_{1,l}, \ldots, a_{n,l}\}$: the set of coordinates recorded in the AIS data for a ship l
$S_l = \{s_{1,l}, \ldots, s_{n,l}\}$: the set of speeds recorded in the AIS data for a ship l
$T_l = \{t_{1,l}, \ldots, t_{n,l}\}$: the set of timestamps recorded in the AIS data for a ship l

Outputs:

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

/ Data Preprocessing */*
$B_l = \{b_{1,l} \ldots b_{k,l}\} \leftarrow$ the set of coordinates in A_l that indicate a speed less than 1 knot
/ Exception Identification */*
if $|B_l| < 100$ **or** $|B_l| > 100,000$ **then**
    Remove the data and stop ▷ *The ship has an abnormal port call*
else
    Continue
end if
/ DBSCAN Clustering */*
$X \leftarrow b_{1,l}$
for $i \leftarrow 2 : k$ **do**
    if $t_i - t_{i-1} \leq \Delta t$ **then**
        Append $b_{i,l}$ to X
    else
        $DBSCAN(X, Eps, MinPts)$
        $X \leftarrow \emptyset$
        Append $b_{i,l}$ to X
    end if
end for
Remove the observations labeled as noise from B_l
Keep only the first observation for each hour in B_l ▷ *Note that only m observations remain in B_l at this stage*
$D_l = \{d_{1,l}, \ldots, d_{m,l}\} \leftarrow$ the coordinates of the first observation for each hour in B_l
$H_l = \{h_{1,l}, \ldots, h_{m,l}\} \leftarrow$ the headings of the first observation for each hour in B_l

Algorithm 2 Level 2 IMA-DBSCAN

Inputs:

$\boldsymbol{D} = \{D_1, \ldots, D_L\}$: the set of coordinates for all ships after Level 1 IMA-DBSCAN
$\boldsymbol{H} = \{H_1, \ldots, H_L\}$: the set of headings for all ships after Level 1 IMA-DBSCAN
$\boldsymbol{O} = \{\boldsymbol{D}, \boldsymbol{H}\} = \{o_1, \ldots, o_M\}$: the combined set of coordinates and headings

Outputs:

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

/ Parameter Initialization */*
$Dist \leftarrow$ the average distance between a point in $\boldsymbol{D}$ and the center of the mass of $\boldsymbol{D}$
$m \leftarrow |\boldsymbol{D}|$
$Eps1 \leftarrow \alpha \cdot Dist$
$MinPts \leftarrow \beta \cdot m$
/ Iteration Process */*
$Dist_0 \leftarrow 0$
while $Dist - Dist_0 > \Delta Dist$ km **do**
    $ST - DBSCAN(\boldsymbol{O}, Eps1, Eps2 = 1°, MinPts)$ ▷ *See function ST-DBSCAN*
    $Dist_0 \leftarrow Dist$
    $Dist \leftarrow$ the average distance between a non-noisy point in $\boldsymbol{D}$ and the center of the mass of its assigned cluster
    $m' \leftarrow$ |noisy points in $\boldsymbol{O}$|
    $NumC \leftarrow$ |clusters in $\boldsymbol{O}$|
    $Eps1 \leftarrow \alpha \cdot Dist$
    $MinPts \leftarrow \beta \cdot (m - m') / NumC$
end while
/ Merging Clusters */*
Use the center of the mass of each cluster to calculate the distance in between
for all clusters c in $\boldsymbol{O}$ **do**
    $c' \leftarrow$ the nearest cluster less than 500 m away from c
    if the docking times of c' and c do not overlap **then**
        Replace the cluster label of c' with that of c
    end if
end for
/ Berth and Anchorage Detection */*
$C_{berth} \leftarrow$ clusters in $\boldsymbol{O}$
$C_{anchorage} \leftarrow DBSCAN$(Noisy points in $\boldsymbol{O}, Eps', MinPts'$)
/ Exception Removal */*
for all clusters c **in** C_{berth} **and** $C_{anchorage}$ **do**
    $NumP \leftarrow$ the number of port calls in cluster c
    if $NumP < N$ **then**
        Remove c
    end if
end for

Algorithm 3 ST-DBSCAN

Inputs:

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

Outputs:

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

```
/* The codes are adapted from those in Birant and Kut (2007).                    */
function ST − DBSCAN(D, Eps1, Eps2, MinPts)
    ClusterLabel = 0
    for i ← 1 : m do
        if o_i is not in a cluster then
            Y ← RetrieveNeighbors(o_i, Eps1, Eps2)        ▷ See function RetrieveNeighbors
            if |Y| < MinPts then
                Mark o_i as noise
            else                                          ▷ Construct a new cluster
                ClusterLabel ← ClusterLabel + 1
                for j ← 1 : |Y| do
                    Mark all objects in Y with current ClusterLabel
                end for
                Push(all objects in Y)
                while not IsEmpty() do
                    CurrentObj = Pop()
                    Z ← RetrieveNeighbors(CurrentObj, Eps1, Eps2)
                    if |Z| ≥ MinPts then
                        for all objects o in Z do
                            if o is not marked as noise or it is not in a cluster then
                                Mark o with current ClusterLabel
                                Push(o)
                            end if
                        end for
                    end if
                end while
            end if
        end if
    end for
    C = {c_1, ..., c_M} ← the set of clusters in O
end function
```

Algorithm 4 RetrieveNeighbors

Inputs:

o: an observation in $\boldsymbol{O}$

$Eps1$: maximum geographical coordinate (spatial) distance

$Eps2$: maximum non-spatial distance

Outputs:

$Neighbors$: the set of neighbors for o

function $RetrieveNeighbors(o, Eps1, Eps2)$
    $Neighbors \leftarrow \emptyset$
    for all observations o' **in** O **do**
        $Dist1 \leftarrow D(o, o')$ ▷ *See Equation* (B.1)
        $Dist2 \leftarrow \Delta h(o, o')$ ▷ *See Equation* (B.2)
        if $Dist1 \leq Eps1$ **and** $Dist2 \leq Eps2$ **then**
            Append o' to $Neighbors$
        end if
    end for
    return $Neighbors$
end function

B.2. Illustrative Case: Port of Ningbo-Zhoushan

To demonstrate the capability of IMA-DBSCAN in accurately identifying anchorage and berth areas of a port, which other methods might fail to achieve, we apply the algorithm to the Port of Ningbo-Zhoushan in China. We choose this specific port due to its intricate layout. Figure B.3a showcases the first 50,000 AIS observations taken since January 1, 2020, within the Port of Ningbo-Zhoushan.[5] The observations are represented by blue dots on the map, with each dot indicating the position of a low-speed container ship. Before applying IMA-DBSCAN to the AIS data, we mark the approximate locations of anchorages and berths using satellite images and nautical charts as benchmarks. The red polygons on the map indicate the anchorage areas, while the yellow rectangles denote the berth locations.

Figure B.3b presents the clustering results of IMA-DBSCAN for the Port of Ningbo-Zhoushan, which mirrors the map in Figure B.3a to allow for a direct comparison between the clustering results and the benchmark observations gathered through satellite images and nautical charts. The clusters in Figure B.3b (colored in red, yellow, blue, purple, cyan, and orange) correspond closely with the anchorage areas in Figure B.3a.[6] Additionally, in Figure B.3e, we spotlight the locations of four terminals within the port of Ningbo-Zhoushan: Beilun, Daxie, Pukou, and Yuandong. Using satellite maps as benchmarks, we confirm the accuracy of these identifications; as shown in the top row of Figure B.4, each berth in the terminals is pinpointed precisely, and the delineated areas align closely with reality.[7]

To assess the performance of IMA-DBSCAN against that of other spatial clustering algorithms, we contrast it with the outcomes from ST-DBSCAN.[8] Given that ST-DBSCAN is capable of processing spatial-temporal databases and is recognized as one of the most prominent spatial clustering algorithms in the literature, this comparison is relevant. Figure B.3c illustrates the results derived from ST-DBSCAN, underscoring its lesser precision in comparison to

[5] Our illustrative case of Ningbo-Zhoushan focuses on a one-month snapshot, as the selected first 50,000 AIS observations spanned throughout January 2020. It is reasonable to assume that the identification results would be indicative of anchorage and berth areas in subsequent months, given that we do not expect significant short-term changes in the port areas. In real-world applications of IMA-DBSCAN, periodic identification can be conducted to monitor potential changes in port anchorages and berths.

[6] For clarity, we also display the convex hulls formed by these clusters in Figure B.3d.

[7] Furthermore, some of the blue dots in Figure B.3a do not correspond to any anchorage or berth in Figure B.3b, indicating that ships anchored in these areas for only a short period.

[8] For this comparison, the input parameters of ST-DBSCAN are set to $Eps1 = 2500$ m, $Eps2 = 1°$, and $MinPts = 100$, as recommended by Ester et al. (1996).

IMA-DBSCAN. Notably, while ST-DBSCAN can generally detect points within the anchorages that are highlighted in blue in Figure B.3c, it mistakenly identifies several high-density regions as berths, even though they are not genuine berths. For example, within the blue rectangle in Figure B.3f, points that ought to be categorized as noise are incorrectly marked as berths, given that ships stayed in these locations for extended periods (potentially for maintenance tasks). Additionally, in the black rectangle, ST-DBSCAN mislabels several points as berths when they should be designated as mooring areas. Consequently, while employing ST-DBSCAN on the sample data offers insights into the arrangement of anchorage areas, it does not succeed in precisely pinpointing berth locations.

Furthermore, in Figure B.4, we present the detailed results of berth identification for each of the four terminals within the Port of Ningbo-Zhoushan: Beilun, Daxie, Pukou, and Yuandong. The outcomes from ST-DBSCAN are ambiguous and feature overlapping sections, which are proximate in position but with significant differences in heading. Although the general range of these terminals can be discerned, individual berths are scarcely distinguishable. In contrast, our IMA-DBSCAN method produces clusters that align precisely with each berth within a terminal. Admittedly, increasing the $MinPts$ value or reducing the $Eps1$ value could enhance the ST-DBSCAN results. However, this would necessitate constant parameter adjustments, which are challenging to execute consistently for each port. Conversely, our IMA-DBSCAN algorithm operates iteratively to automatically determine a set of parameters that can accurately identify both berths and anchorages.

Lastly, we apply the same exercise to the Ports of Los Angeles and Long Beach in the U.S., Rotterdam in the Netherlands, and Singapore. As shown in Figures B.5 through B.7, we confirm that our IMA-DBSCAN algorithm delivers more accurate identification results than the traditional ST-DBSCAN algorithm across all these major container ports.

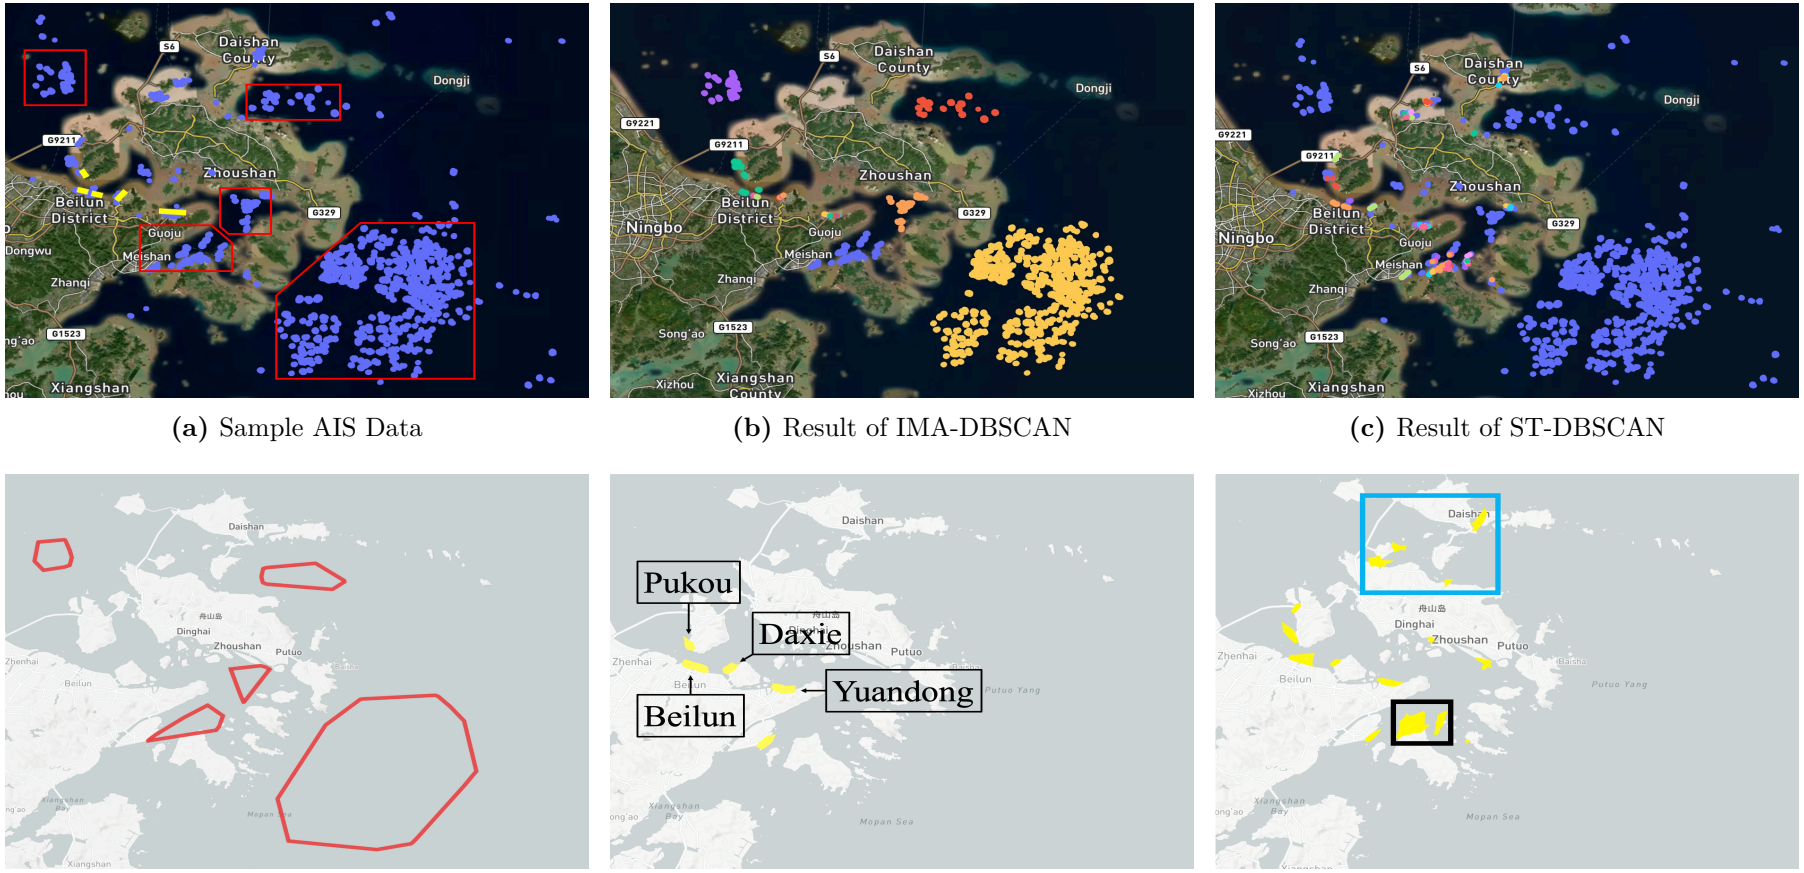

(a) Sample AIS Data **(b)** Result of IMA-DBSCAN **(c)** Result of ST-DBSCAN

(d) Anchorages (IMA-DBSCAN) **(e)** Berths (IMA-DBSCAN) **(f)** Berths (ST-DBSCAN)

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

Notes. In Figure (a), the sample data comprise the first 50,000 AIS observations taken since January 1, 2020, within the Port of Ningbo-Zhoushan in China. These observations are represented by blue dots on the map, corresponding to coordinates ranging from 121.60°E to 123.00°E and from 29.50°N to 30.35°N. As a benchmark, using satellite maps and nautical charts, we identify the approximate areas of the anchorages with red polygons and the approximate locations of the berths with yellow rectangles. We apply two clustering algorithms, IMA-DBSCAN and ST-DBSCAN, to the sample data. The resulting clusters are depicted in Figures (b) and (c) respectively. Notably, the blue dots in Figure (b) represent an identified anchorage area, while those in Figure (c) represent noise, which outlines the general layout of anchorage areas but does not distinctly identify each one. In Figure (d), the anchorages from Figure (b) are shown separately in red polygons. In Figure (e), the berths from Figure (b) are displayed separately in yellow. The four terminals are identified as Pukou, Daxie, Beilun, and Yuandong. Lastly, in Figure (f), the yellow areas depict the approximate positions of the berths as identified by ST-DBSCAN, while the blue and black rectangles indicate misidentifications of noise as berths and confusion between anchorages and berths, respectively.

(a) Terminal Beilun **(b)** Terminal Daxie **(c)** Terminal Pukou **(d)** Terminal Yuandong

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

Notes. Detailed results of berth identification for each of the four terminals: Beilun, Daxie, Pukou, and Yuandong, within the Port of Ningbo-Zhoushan. The berths identified by IMA-DBSCAN are presented in yellow on the top row, while those pinpointed by ST-DBSCAN are depicted in brown on the bottom row.

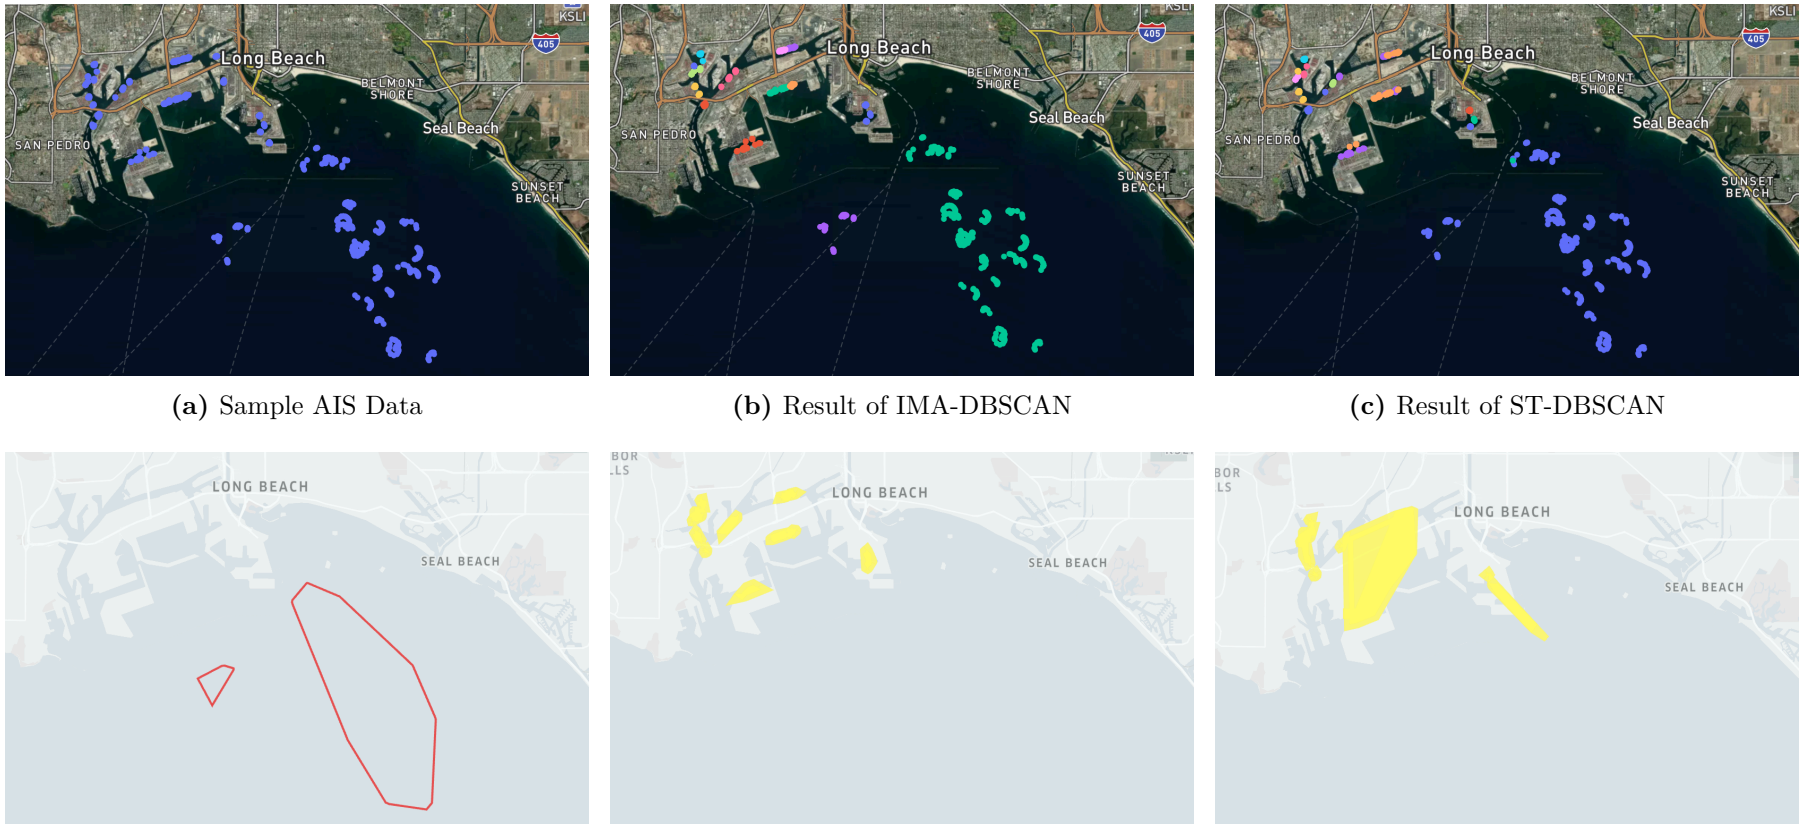

(a) Sample AIS Data

(b) Result of IMA-DBSCAN

(c) Result of ST-DBSCAN

(d) Anchorages (IMA-DBSCAN)

(e) Berths (IMA-DBSCAN)

(f) Berths (ST-DBSCAN)

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

Notes. In Figure (a), the sample data comprise the first 50,000 AIS observations taken since January 1, 2020, within the Ports of Los Angeles and Long Beach in the U.S., represented by blue dots on the map. We apply two clustering algorithms, IMA-DBSCAN and ST-DBSCAN, to the sample data. The resulting clusters are depicted in Figures (b) and (c), respectively. Notably, the blue dots in Figure (b) represent an identified anchorage area, while those in Figure (c) represent noise, which outlines the general layout of anchorage areas but does not distinctly identify each one. In Figure (d), the anchorages from Figure (b) are shown separately in red polygons. In Figure (e), the berths from Figure (b) are displayed separately in yellow. Lastly, in Figure (f), the yellow areas depict the approximate positions of the berths as identified by ST-DBSCAN.

(a) Sample AIS Data **(b)** Result of IMA-DBSCAN **(c)** Result of ST-DBSCAN

(d) Anchorages (IMA-DBSCAN) **(e)** Berths (IMA-DBSCAN) **(f)** Berths (ST-DBSCAN)

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

Notes. In Figure (a), the sample data comprise the first 50,000 AIS observations taken since January 1, 2020, within the Port of Rotterdam in the Netherlands, represented by blue dots on the map. We apply two clustering algorithms, IMA-DBSCAN and ST-DBSCAN, to the sample data. The resulting clusters are depicted in Figures (b) and (c), respectively. Notably, the blue dots in Figure (b) represent an identified anchorage area, while those in Figure (c) represent noise, which outlines the general layout of anchorage areas but does not distinctly identify each one. In Figure (d), the anchorages from Figure (b) are shown separately in red polygons. In Figure (e), the berths from Figure (b) are displayed separately in yellow. Lastly, in Figure (f), the yellow areas depict the approximate positions of the berths as identified by ST-DBSCAN.

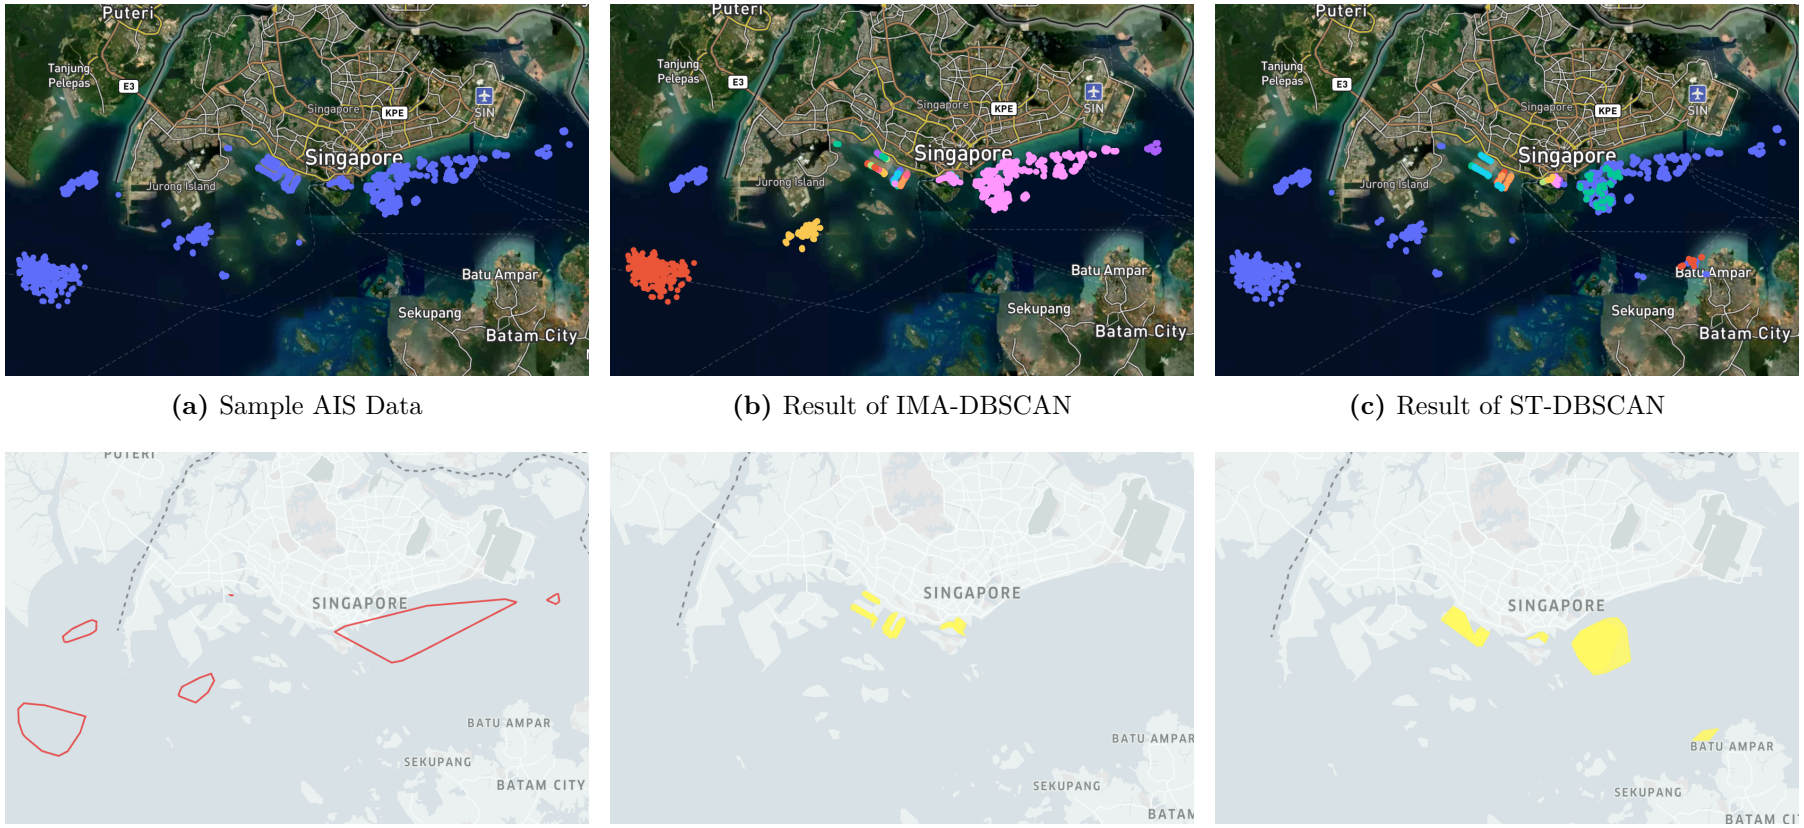

(a) Sample AIS Data **(b)** Result of IMA-DBSCAN **(c)** Result of ST-DBSCAN

(d) Anchorages (IMA-DBSCAN) **(e)** Berths (IMA-DBSCAN) **(f)** Berths (ST-DBSCAN)

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

Notes. In Figure (a), the sample data comprise the first 50,000 AIS observations taken since January 1, 2020, within the Port of Singapore, represented by blue dots on the map. We apply two clustering algorithms, IMA-DBSCAN and ST-DBSCAN, to the sample data. The resulting clusters are depicted in Figures (b) and (c), respectively. Notably, the blue dots in Figure (b) represent an identified anchorage area, while those in Figure (c) represent noise, which outlines the general layout of anchorage areas but does not distinctly identify each one. In Figure (d), the anchorages from Figure (b) are shown separately in red polygons. In Figure (e), the berths from Figure (b) are displayed separately in yellow. Lastly, in Figure (f), the yellow areas depict the approximate positions of the berths as identified by ST-DBSCAN.

B.3. Weekly Indices of Port Congestion

The integration of high-frequency AIS data with our IMA-DBSCAN algorithm enables the construction of port congestion indices at even higher frequencies than monthly updates. The AIS system processes over 2,000 reports per minute, with the capacity to update information as frequently as every two seconds. This remarkable rate of data collection allows for the capture of detailed, real-time movements of ocean-going vessels. Furthermore, unlike many traditional algorithms that require data at specific time intervals to function effectively, IMA-DBSCAN is uniquely flexible. Its streamlined design allows it to operate independently of fixed time frequencies, making it particularly suitable for working with the highly variable and high-frequency data provided by AIS.

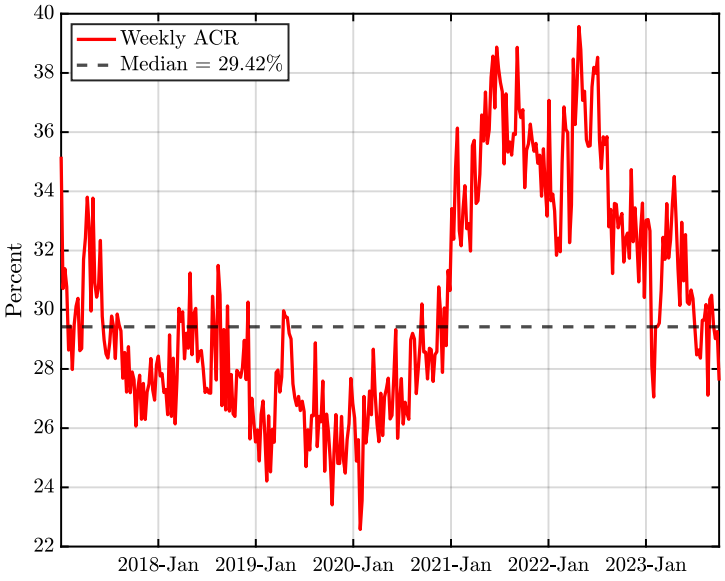

Figure B.8: Weekly ACR Index

Notes. The weekly ACR index is derived in the same way as the monthly ACR index, except that the construction process – i.e., counting the number of delayed ship visits, normalizing using the total number of ship visits, and aggregating into a time series measure – is now based on weekly updates of the AIS data. The index is presented in percentage terms.

Figure B.8 presents our ACR index at a weekly frequency. Compared to the monthly ACR index in Figure 4, we can see that the general trends between the two series are essentially the same. However, the weekly series is significantly more volatile than the monthly series, as the number of ship visits to each port might vary considerably during each week.[9]

[9]Technically, we could also construct the ACR index based on daily or even hourly updates of the AIS data. However, since there are only a few ship visits to a port each day or hour, the constructed ACR index would

In addition to the ACR index, we also introduce an alternative high-frequency congestion metric for container ports, namely, the Average Congestion Time (ACT). Unlike the ACR index, the ACT index measures the average number of hours a container ship waits in an anchorage area of a port before docking at a berth, weighted by the relative number of ship visits to each of the top 50 container ports worldwide:

$$ACT_t = \sum_{p \in \mathscr{P}} \left[\frac{Delayed_{p,t} + Undelayed_{p,t}}{\sum_{p \in \mathscr{P}} (Delayed_{p,t} + Undelayed_{p,t})} \cdot \frac{DelayHours_{p,t}}{Delayed_{p,t} + Undelayed_{p,t}} \right],$$

where $Delayed_{p,t}$, $Undelayed_{p,t}$, and $DelayHours_{p,t}$ represent the number of delayed and undelayed ship visits, and the total number of hours that container ships spend in the anchorage areas of port p in week t, respectively. Figure B.9 plots the ACT index at a weekly frequency, which is observed to closely co-move with the ACR index during the sample period. In Appendix F.6, we also reduce the time dimension of our ACT index to a monthly frequency and use the resulting series in the causality assessment. We show that our identification results are robust when using the ACT index as the measure indicative of the state of the global supply chain.

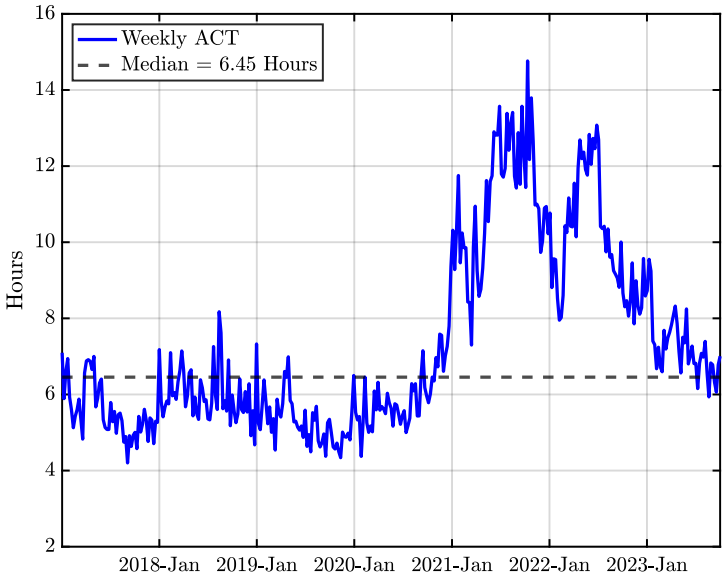

Figure B.9: Weekly ACT Index

Notes. The ACT index is derived by taking a weighted average of the number of hours a container ship waits in an anchorage area of a port before docking at a berth, with the weights based on the relative number of ship visits to each port.

become extremely volatile, with 0% and 100% occurring frequently.

C. Discussion on the Assumptions in the Model

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

The matching function. There is ample literature studying the origins of matching frictions in the product market, including but not limited to locating and building connections with buyers in different locations (Benguria, 2021; Krolikowski and McCallum, 2021; Lenoir et al., 2022), the costly acquisition of information about market conditions elsewhere (Allen, 2014; Chaney, 2014), and informal trade barriers such as common language (Melitz and Toubal, 2014) and geography (Eaton and Kortum, 2002). Common across all these theories is the presence of prevalent trade barriers between producers and retailers, implying that not all unmatched producers engage in trade and not all sourcing visits by retailers are successful. Consistent with these theories, we assume a constant-returns-to-scale matching function that summarizes how unmatched producers and visits by retailers are transformed into trades through the matching process. Thus, we abstract from modeling the complex matching process while still preserving its main implication: the unmatched producers only engage in trade with probability $f(\theta)G(\tau) < 1$, and the visits of retailers to producers are only successful with probability $q(\theta)G(\tau) < 1$.

Endogenous separation on transportation cost. In much the same way as the separation margin in the labor market can be modeled endogenously when workers face productivity shocks to their employment matches, and bad draws possibly lead to separations (Bils et al., 2011; Menzio and Shi, 2011; Fujita and Ramey, 2012), the separation margin in the product market can also be modeled endogenously. This occurs when producers face idiosyncratic transportation costs to their trading relationships with retailers, and bad draws possibly lead to the termination of such relationships. Such a modeling assumption is reasonable only if we find convincing evidence that

(i) transportation costs are taken into account when trading partners decide on a potential trade, and (ii) there exists a threshold of transportation costs above which trading partners choose to sever their relationship.

The prediction that transportation cost affects the probability of trade has been empirically examined in the trade literature. To name a few examples, evidence in Rodrigue (2020) underlines that across all modes, increasing transportation costs by 10% reduces trade volumes by more than 20%. In the context of maritime transportation, Brancaccio et al. (2020) exploit changes in tariffs across the trade network to estimate the elasticity of world trade value with respect to shipping costs. They estimate that a 1% change in shipping costs leads to approximately a 1% change in world trade value. Similarly, Wong (2022) estimates the containerized trade elasticity with respect to freight rates using the round-trip effect as an instrument. In particular, for route i, j, she uses a Bartik-style instrument to proxy for the predicted trade value on route j, i. Wong (2022) reports that a 1% increase in per-unit container freight rates decreases containerized trade value by 2.8% when dyad-by-product controls are included in the regression. Together, these estimated elasticities emphasize that transportation costs are indeed taken into account when trading partners choose to form a relationship and that a rise in transportation costs deters trade substantially.

Both theory and casual observation also suggest the presence of a reservation transportation cost, which trading partners often consider when assessing the profitability of potential trades in the product market. Notably, to reconcile the empirical evidence that export and import intensities vary across plants, Kasahara and Lapham (2013) extend the model in Melitz (2003) by allowing for heterogeneity in transportation costs. Incorporating heterogeneous transportation costs provides a plausible self-selection mechanism for trade decisions, as plants with low transportation costs self-select into exporting and importing. Consequently, there must exist a threshold value of transportation cost below which plants choose to engage in trade.[10] Furthermore, as in our equilibrium selection mechanism, the reservation transportation cost is likely to be largely fixed in the short run, since it represents the minimum transportation expense necessary

[10] A similar argument is found in the discussion of transport infrastructure and its effects on a firm's exporting decision. For instance, Naudé and Matthee (2011) argue that the availability of transport infrastructure will have a threshold effect – a certain minimum level of infrastructure is required for a firm to start exporting. However, once this threshold is reached, improved infrastructure will not necessarily have a large impact on the extent of an individual firm's exports. Since the availability of transport infrastructure (at least partially) determines the transportation cost, this argument is consistent with the presence of a reservation transportation cost that firms take into account when making exporting decisions.

for engaging in trade. This cost includes fixed elements such as fuel costs, labor, insurance, and tariffs, which are not easily alterable in the short term. Economically, it is reasonable to perceive this cost as fixed because trading partners often negotiate these expenses based on long-term contracts or established trade routes, which are not frequently revised. For instance, Walmart has established long-term shipping contracts with Maersk Line to secure stable transportation costs for moving goods from Asia to North America. These contracts, often set for multiple years, lock in shipping rates and other associated costs, ensuring Walmart can anticipate its logistics expenses despite market fluctuations.

D. Long Proofs and Model Dynamics

The following appendices include all long proofs that are omitted from the main text, as well as several important discussions on model dynamics. For simplicity, in the proofs presented from Appendix D.1 to D.8, we omit the matching efficiency parameter A, since it is parameterized to one. In Appendix D.9, where the supply chain disruption is modeled as a reduction in matching efficiency, and in Appendix D.10, where we examine the convergence dynamics of the model from one steady state to another following each shock of interest, A is explicitly incorporated in all derivations.

D.1. Proof of Proposition 1

We first rewrite $\mathbb{E}_{z'}S(z')$. Using the definition of $S(z)$ in Equation (10), we derive the following expression for the total surplus:

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

Subtracting Equation (14) from the above equation yields $S(z) = (\bar{z} - z)l$. Replacing $S(z')$ in $\mathbb{E}_{z'}S(z')$ with $(\bar{z} - z')l$, we derive that:

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

Subsequently, by replacing $\mathbb{E}_{z'}S(z')$ in Equations (14) and (15) with $l\int_0^{\bar{z}} G(z')dz'$, the match separation condition can be re-written as:

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

and the match creation condition is given by:

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

Combining the two conditions yields Equation (16):

$$
\begin{aligned}
0 &= p - \bar{z} + (1 - \eta f(\theta))\,\beta \int_0^{\bar{z}} G(z')dz' + \frac{\rho}{q(\theta)} - (1-\eta)\beta \int_0^{\bar{z}} G(z')dz' \\
0 &= p - \bar{z} - \eta f(\theta)\beta \int_0^{\bar{z}} G(z')dz' + \frac{\rho}{q(\theta)} + \eta\beta \int_0^{\bar{z}} G(z')dz' \\
0 &= p - \bar{z} - \eta f(\theta)\frac{\rho}{q(\theta)(1-\eta)} + \beta \int_0^{\bar{z}} G(z')dz' \\
\theta &= \frac{1-\eta}{\eta\rho}\left(p - \bar{z} + \beta \int_0^{\bar{z}} G(z')dz'\right),
\end{aligned}
$$

where the last step is obtained using the property $\theta = f(\theta)/q(\theta)$.

The first property is obvious. Since θ cannot be negative, for a given $\bar{z}$, p is bounded on $[p^{min}, +\infty)$, where p^{min} is such that it solves $p^{min} - \bar{z} + \beta \int_0^{\bar{z}} G(z')dz' = 0$ for any $\bar{z} > 0$. As for the second and third properties, we get that:

$$
\frac{\partial \theta(p, \bar{z})}{\partial p} = \frac{1-\eta}{\eta\rho} > 0.
$$

Therefore, the equilibrium product market tightness $\theta(p, \bar{z})$ is strictly increasing and linear on $[p^{min}, +\infty)$.

D.2. Proof of Proposition 2

The first property is obvious. When $p = p^{min}$, we have $\theta(p^{min}) = 0$, $f\left(\theta(p^{min})\right) = 0$, and $c_s^{flex}(p^{min}) = 0$. When $p \to +\infty$, since A is parameterized to one, we have $\lim_{p\to+\infty} \theta(p) = +\infty$, $\lim_{p\to+\infty} f\left(\theta(p)\right) = 1$, and hence $\lim_{p\to+\infty} c_s^{flex}(p) = G(\tau)l$. In terms of the second and third properties, we derive that:

$$
\frac{dc_s^{flex}(p)}{dp} = \frac{1-\eta}{\eta\rho} \frac{\left(1 - G(\tau)\right) q(\theta)^{1+\xi} G(\tau) l}{\left(1 - G(\tau) + f(\theta)G(\tau)\right)^2} > 0,
$$

$$
\begin{aligned}
\frac{d^2 c_s^{flex}(p)}{dp^2} = &- \left(\frac{1-\eta}{\eta\rho}\right)^2 \left(1 - G(\tau)\right) G(\tau) l \\
&\cdot \frac{\left(1 - G(\tau) + f(\theta)G(\tau)\right) \theta^{\xi-1} (1+\xi)(1+\theta^{\xi})^{-\frac{1+\xi}{\xi}-1} + 2G(\tau) q(\theta)^{2(1+\xi)}}{\left(1 - G(\tau) + f(\theta)G(\tau)\right)^3} < 0.
\end{aligned}
$$

Consequently, the flexible price aggregate supply c_s^{flex} is strictly increasing and concave on $[p^{min}, +\infty)$.

D.3. Proof of Proposition 4

Since we look for a flexible price steady state with positive consumption, we restrict our search of price p within the range $[p^{min}, +\infty)$. The steady state condition (26) can be re-written as the following:

$$\frac{f\left(\theta(p)\right)p}{1-G(\tau)+f\left(\theta(p)\right)G(\tau)} = \chi^{\varepsilon}\frac{\mu}{G(\tau)l}. \tag{D.1}$$

For any $\tau > 0$, the right-hand side is a strictly positive constant. For the left-hand side, when $p = p^{min}$, $\theta(p^{min}) = 0$, and $f\left(\theta(p^{min})\right) = 0$, it has a limit of zero; when $p \to +\infty$, $\lim_{p\to+\infty}\theta(p) = +\infty$, and $\lim_{p\to+\infty} f\left(\theta(p)\right) = 1$, it has a limit of positive infinity. For $p \in [p^{min}, +\infty)$, the derivative of the left-hand side with respect to p is given by:

$$\frac{d}{dp}\left[\frac{f\left(\theta(p)\right)p}{1-G(\tau)+f\left(\theta(p)\right)G(\tau)}\right] = \frac{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}{\left(1-G(\tau)+f(\theta)G(\tau)\right)^2} > 0.$$

Therefore, the left-hand side is strictly increasing from zero to positive infinity on $[p^{min}, +\infty)$ and there is a unique $p \in [p^{min}, +\infty)$ that solves Equation (D.1).

D.4. Proof of Proposition 5

We first consider an adverse shock to aggregate demand. No matter whether the negative aggregate demand shock is represented by a decrease in the money supply μ or in the taste for consumption of goods χ, the right-hand side of Equation (D.1) will decrease. To balance both sides of Equation (D.1), price p will decrease since the derivative of the left-hand side with respect to p is positive (see the proof for Proposition 4). As p decreases, by the second property in Proposition 1 and the second property in Proposition 2, product market tightness θ and consumption (or, equivalently, output) c will decrease. Since both p and θ decrease, according to Equation (12), wholesale price r will decline as well. For the matching cost, $G(\tau)l - c$, and spare capacity, $l - c$, since they both are strictly decreasing in c, they will increase following an adverse shock to aggregate demand.

Next, we consider an adverse shock to productive capacity, which is parameterized by a decrease in the fixed factor endowment of producers l. On impact, the right-hand side of Equation (D.1) increases. To balance both sides of Equation (D.1), price p will increase since the derivative of the left-hand side with respect to p is positive. As p increases, by the second property in

Proposition 1 and Proposition 3, product market tightness θ will increase while consumption (or, equivalently, output) c will fall. Again, since both p and θ increase, according to Equation (12), wholesale price r will rise. In terms of the matching cost and spare capacity, they can be alternatively expressed as the following:

$$matching\ cost = G(\tau)l - c = \frac{(1-G(\tau))\,(1-f(\theta))}{1-(1-f(\theta))\,G(\tau)}G(\tau)l,$$

$$spare\ capacity = l - c = \frac{1-G(\tau)}{1-G(\tau)+f(\theta)G(\tau)}l,$$

respectively. With θ increasing and l decreasing following the negative shock to productive capacity, it is straightforward to verify that both the matching cost and spare capacity will decrease.

Lastly, we consider an adverse shock to the supply chain, which is represented by an increase in γ, i.e., the scale parameter of the log-normal distribution of transportation costs $G(.)$. We first look at the effect on price p. Using the re-arranged steady state condition (D.1), we define a function $\mathbb{T} : [p^{min}, +\infty) \times \mathbb{R} \to \mathbb{R}$:

$$\begin{aligned}\mathbb{T}(p,\gamma) &= \chi^{\varepsilon}\frac{\mu}{G(\tau)l} - \frac{f(\theta)p}{1-G(\tau)+f(\theta)G(\tau)} \\ &= \chi^{\varepsilon}\frac{\mu}{\Phi(\frac{\log\tau-\gamma}{\sigma})l} - \frac{\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^{\tau}\Phi(\frac{\log z'-\gamma}{\sigma})dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}p}{1-\Phi(\frac{\log\tau-\gamma}{\sigma})+\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^{\tau}\Phi(\frac{\log z'-\gamma}{\sigma})dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}\Phi(\frac{\log\tau-\gamma}{\sigma})},\end{aligned}$$

where $\Phi(.)$ is the standard normal cumulative density function. Assuming the existence of a tuple $(p_0, \gamma_0) \in [p^{min}, +\infty) \times \mathbb{R}$ such that $\mathbb{T}(p_0, \gamma_0) = 0$ and $\partial\mathbb{T}(p,\gamma)/\partial p|_{p=p_0,\gamma=\gamma_0} \neq 0$, by the implicit function theorem, there is a neighborhood of (p_0, γ_0) such that whenever γ is sufficiently close to γ_0, there is a unique p so that $\mathbb{T}(p, \gamma) = 0$. This assignment makes p a continuous function of γ. Applying implicit differentiation to $\mathbb{T}(p, \gamma)$ around (p_0, γ_0) yields:

$$\frac{dp(\gamma)}{d\gamma} = -\frac{\partial\mathbb{T}(p,\gamma)/\partial\gamma}{\partial\mathbb{T}(p,\gamma)/\partial p}.$$

In terms of $\partial\mathbb{T}(p,\gamma)/\partial\gamma$, we derive that:

$$\begin{aligned}\frac{\partial\mathbb{T}(p,\gamma)}{\partial\gamma} =& \chi^{\varepsilon}\frac{\mu}{l}\frac{\frac{1}{\sigma}g(\tau)}{G(\tau)^2} + \frac{(1-G(\tau))\,\frac{(1-\eta)\beta}{\eta\rho}\left[\int_0^{\tau}\frac{1}{\sigma}g(z')dz'\right]q(\theta)^{1+\xi}p}{(1-G(\tau)+f(\theta)G(\tau))^2} \\ &+ \frac{(1-f(\theta))\,f(\theta)p\frac{1}{\sigma}g(\tau)}{(1-G(\tau)+f(\theta)G(\tau))^2} > 0,\end{aligned}$$

where $g(\tau) \equiv \phi[(\log\tau - \gamma)/\sigma]$, $g(z') \equiv \phi[(\log z' - \gamma)/\sigma]$, while $\phi(.)$ is the standard normal probability density function. In terms of $\partial\mathbb{T}(p,\gamma)/\partial p$, it can be written as follows:

$$\frac{\partial\mathbb{T}(p,\gamma)}{\partial p} = -\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}{(1-G(\tau)+f(\theta)G(\tau))^2} < 0.$$

By combining $\partial\mathbb{T}(p,\gamma)/\partial\gamma$ with $\partial\mathbb{T}(p,\gamma)/\partial p$ and collecting terms, we pin down the derivative of p with respect to γ:

$$\begin{aligned}\frac{dp(\gamma)}{d\gamma} = &\left[(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))\right]^{-1}\\ &\cdot\Big\{\,(1-G(\tau)+f(\theta)G(\tau))\,f(\theta)\frac{1}{\sigma}g(\tau)\frac{p}{G(\tau)}\\ &+(1-G(\tau))\,\frac{(1-\eta)\beta}{\eta\rho}\left[\int_0^\tau\frac{1}{\sigma}g(z')dz'\right]q(\theta)^{1+\xi}p\\ &+(1-f(\theta))\,f(\theta)p\frac{1}{\sigma}g(\tau)\Big\} > 0.\end{aligned} \tag{D.2}$$

Hence, price p will increase on impact after an adverse shock to the supply chain. In terms of consumption (or, equivalently, output), it is written as follows:

$$c(\gamma) = \chi^{\varepsilon}\frac{\mu}{p(\gamma)},$$

where p is an implicit function of γ. Therefore, the derivative of c with respect to γ is given by the following:

$$\begin{aligned}\frac{dc(\gamma)}{d\gamma} = &-\chi^{\varepsilon}\frac{\mu}{p}\left[(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))\right]^{-1}\\ &\cdot\Big\{\,(1-G(\tau)+f(\theta)G(\tau))\,f(\theta)\frac{1}{\sigma}g(\tau)\frac{1}{G(\tau)}\\ &+(1-G(\tau))\,\frac{(1-\eta)\beta}{\eta\rho}\left[\int_0^\tau\frac{1}{\sigma}g(z')dz'\right]q(\theta)^{1+\xi}\\ &+(1-f(\theta))\,f(\theta)\frac{1}{\sigma}g(\tau)\Big\} < 0.\end{aligned}$$

Hence, consumption (or, equivalently, output) c will fall. Next, in terms of product market tightness, it is given by:

$$\theta(\gamma) = \frac{1-\eta}{\eta\rho}\left(p(\gamma) - \tau + \beta\int_0^\tau\Phi(\frac{\log z' - \gamma}{\sigma})dz'\right). \tag{D.3}$$

Accordingly, the derivative of θ with respect to γ is:

$$\begin{aligned}\frac{d\theta(\gamma)}{d\gamma} =& \frac{1-\eta}{\eta\rho}\left[(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))\right]^{-1} \\ &\cdot\Big\{\,(1-G(\tau)+f(\theta)G(\tau))\,f(\theta)\frac{1}{\sigma}g(\tau)\frac{p}{G(\tau)} + (1-f(\theta))\,f(\theta)p\frac{1}{\sigma}g(\tau) \\ &-(1-G(\tau)+f(\theta)G(\tau))\,f(\theta)\beta\left[\int_0^\tau \frac{1}{\sigma}g(z')dz'\right]\Big\},\end{aligned} \tag{D.4}$$

whose value depends on the values of θ and p. As we will discuss later in Appendix D.6, this dependence is crucial for our discussion on the effectiveness of a contractionary monetary policy shock in controlling inflation and output. Similarly, by substituting Equation (D.3) into Equation (12), we can express the wholesale price as follows:

$$r(\gamma) = p(\gamma) + (1-\eta)\beta\int_0^\tau \Phi(\frac{\log z' - \gamma}{\sigma})dz' + (1-\eta)(z-\tau).$$

Differentiating $r(\gamma)$ with respect to γ yields:

$$\begin{aligned}\frac{dr(\gamma)}{d\gamma} =& \left[(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))\right]^{-1} \\ &\cdot\Big\{\,(1-G(\tau)+f(\theta)G(\tau))\,f(\theta)\frac{1}{\sigma}g(\tau)\frac{p}{G(\tau)} \\ &+(1-G(\tau))\,\frac{(1-\eta)\beta}{\rho}\left[\int_0^\tau \frac{1}{\sigma}g(z')dz'\right]q(\theta)^{1+\xi}p \\ &+(1-f(\theta))\,f(\theta)p\frac{1}{\sigma}g(\tau) \\ &-(1-G(\tau)+f(\theta)G(\tau))\,f(\theta)(1-\eta)\beta\left[\int_0^\tau \frac{1}{\sigma}g(z')dz'\right]\Big\},\end{aligned}$$

whose value is also dependent on the values of θ and p. As for the matching cost, given that it is measured by the difference between $G(\tau)l$ and c, its derivative with respect to γ can be written as follows:

$$\begin{aligned}\frac{d}{d\gamma}\left[matching\ cost(\gamma)\right] =& \left[(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))\right]^{-1} \\ &\cdot\Big(\chi^{\varepsilon}\frac{\mu}{p}\Big\{\,(1-G(\tau)+f(\theta)G(\tau))\,f(\theta)\frac{1}{\sigma}g(\tau)\frac{1}{G(\tau)} \\ &+(1-G(\tau))\,\frac{(1-\eta)\beta}{\eta\rho}\left[\int_0^\tau \frac{1}{\sigma}g(z')dz'\right]q(\theta)^{1+\xi} \\ &+(1-f(\theta))\,f(\theta)\frac{1}{\sigma}g(\tau)\Big\}\end{aligned}$$

$$
\begin{aligned}
&- (1 - G(\tau)) \frac{1-\eta}{\eta\rho} q(\theta)^{1+\xi} p \frac{1}{\sigma} g(\tau) l \\
&- f(\theta) (1 - G(\tau) + f(\theta) G(\tau)) \frac{1}{\sigma} g(\tau) l \Big),
\end{aligned}
$$

whose value is again dependent on the values of θ and p. On the contrary, since the spare capacity is measured by the difference between l and c, its derivative with respect to γ is positive, i.e.,

$$
\begin{aligned}
\frac{d}{d\gamma} \left[spare\ capacity(\gamma) \right] =& \chi^{\varepsilon} \frac{\mu}{p} \left[(1 - G(\tau)) \frac{1-\eta}{\eta\rho} q(\theta)^{1+\xi} p + f(\theta) (1 - G(\tau) + f(\theta) G(\tau)) \right]^{-1} \\
& \cdot \Big\{ (1 - G(\tau) + f(\theta) G(\tau)) f(\theta) \frac{1}{\sigma} g(\tau) \frac{1}{G(\tau)} \\
& + (1 - G(\tau)) \frac{(1-\eta)\beta}{\eta\rho} \left[\int_0^{\tau} \frac{1}{\sigma} g(z') dz' \right] q(\theta)^{1+\xi} \\
& + (1 - f(\theta)) f(\theta) \frac{1}{\sigma} g(\tau) \Big\} > 0.
\end{aligned}
$$

D.5. Slope of the Aggregate Supply Curve and Its Dependence on Product Market Tightness

Given Equation (22), the slope of the aggregate supply curve is determined by the following:

$$
\frac{dc_s^{flex}}{dp} = \frac{(1-\eta) l}{\eta\rho} \frac{(1 - G(\tau)) G(\tau) \left(1 + \theta^{\xi}\right)^{-\frac{1+\xi}{\xi}}}{\left[1 - G(\tau) + \left(1 + \theta^{-\xi}\right)^{-\frac{1}{\xi}} G(\tau)\right]^2}.
$$

Differentiating it with respect to product market tightness θ yields:

$$
\begin{aligned}
\frac{d^2 c_s^{flex}}{dp d\theta} =& - \frac{(1-\eta) l}{\eta\rho} \frac{1}{\left[1 - G(\tau) + \left(1 + \theta^{-\xi}\right)^{-\frac{1}{\xi}} G(\tau)\right]^3} \\
& \cdot \Big\{ \left[1 - G(\tau) + \left(1 + \theta^{-\xi}\right)^{-\frac{1}{\xi}} G(\tau)\right] (1 - G(\tau)) G(\tau) \theta^{\xi - 1} (1 + \xi) \left(1 + \theta^{\xi}\right)^{-\frac{1+\xi}{\xi} - 1} \\
& + 2 (1 - G(\tau)) G(\tau)^2 \left(1 + \theta^{\xi}\right)^{-\frac{2(1+\xi)}{\xi}} \Big\} < 0.
\end{aligned}
$$

Therefore, the aggregate supply curve becomes steeper in the (Q, p) plane as product market tightness increases.

Furthermore, to complement our discussion in Section 3.4 of the main text regarding the response of the system to an adverse supply chain disturbance, we depict an alternative scenario in which the resulting price increase fails to elevate product market tightness at the new steady state in the following figure.

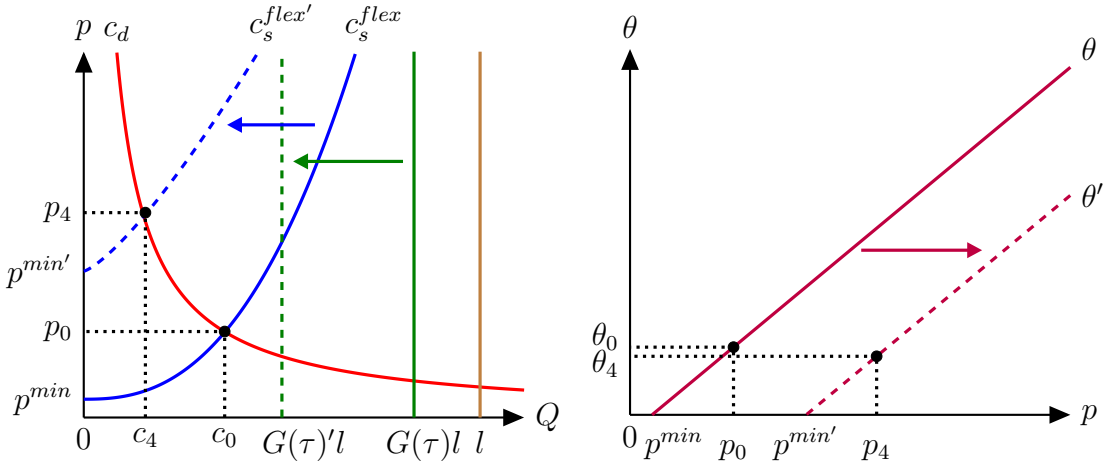

Figure D.1: Alternative Scenario When the Resulting Price Increase Fails to Elevate Product Market Tightness After a Supply Chain Disturbance

D.6. Proof of Proposition 6

We start the proof by re-visiting the function $\mathbb{T} : [p^{min}, +\infty) \times \mathbb{R}^+ \times \mathbb{R} \to \mathbb{R}$:

$$\mathbb{T}(p,\mu,\gamma) = \chi^{\varepsilon}\frac{\mu}{G(\tau)l} - \frac{f(\theta)p}{1-G(\tau)+f(\theta)G(\tau)}$$
$$= \chi^{\varepsilon}\frac{\mu}{\Phi(\frac{\log\tau-\gamma}{\sigma})l} - \frac{\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^{\tau}\Phi(\frac{\log z'-\gamma}{\sigma})dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}p}{1-\Phi(\frac{\log\tau-\gamma}{\sigma})+\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^{\tau}\Phi(\frac{\log z'-\gamma}{\sigma})dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}\Phi(\frac{\log\tau-\gamma}{\sigma})},$$

where $\Phi(.)$ is the standard normal cumulative density function. Assuming the existence of a tuple $(p_0,\mu_0,\gamma_0) \in [p^{min},+\infty)\times\mathbb{R}^+\times\mathbb{R}$ such that $\mathbb{T}(p_0,\mu_0,\gamma_0)=0$ and $\partial\mathbb{T}(p,\mu,\gamma)/\partial p|_{p=p_0,\mu=\mu_0,\gamma=\gamma_0} \neq 0$, by the implicit function theorem, there is a neighborhood of (p_0,μ_0,γ_0) such that whenever (μ,γ) is sufficiently close to (μ_0,γ_0), there is a unique p so that $\mathbb{T}(p,\mu,\gamma)=0$. This assignment makes p a continuous function of μ and γ. Applying implicit differentiation to $\mathbb{T}(p,\mu,\gamma)$ around (p_0,μ_0,γ_0) yields:

$$\frac{\partial p(\mu,\gamma)}{\partial\mu} = -\frac{\partial\mathbb{T}(p,\mu,\gamma)/\partial\mu}{\partial\mathbb{T}(p,\mu,\gamma)/\partial p}.$$

The numerator can be written as:

$$\frac{\partial\mathbb{T}(p,\mu,\gamma)}{\partial\mu} = \chi^{\varepsilon}\frac{1}{G(\tau)l} > 0,$$

whereas the denominator is given by:

$$\frac{\partial \mathbb{T}(p,\mu,\gamma)}{\partial p} = -\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}{(1-G(\tau)+f(\theta)G(\tau))^2} < 0.$$

By combining $\partial\mathbb{T}(p,\mu,\gamma)/\partial\mu$ with $\partial\mathbb{T}(p,\mu,\gamma)/\partial p$, we derive that:

$$\begin{aligned}\frac{\partial p(\mu,\gamma)}{\partial \mu} &= \chi^{\varepsilon}\frac{1}{G(\tau)l}\frac{(1-G(\tau)+f(\theta)G(\tau))^2}{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}\\ &= \frac{1}{\mu}\left[\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))} + \frac{1}{p}\right]^{-1} > 0,\end{aligned}$$

where the last step is obtained because:

$$\mathbb{T}(p,\mu,\gamma) = 0 \;\Rightarrow\; \chi^{\varepsilon}\frac{1}{G(\tau)l} = \frac{f(\theta)p}{\mu\,(1-G(\tau)+f(\theta)G(\tau))}.$$

In terms of the partial derivative of consumption (or, equivalently, output) with respect to μ, using the expression for the aggregate demand in Equation (25), we can derive that:

$$\begin{aligned}\frac{\partial c(\mu,\gamma)}{\partial \mu} &= \chi^{\varepsilon}\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p + f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}\\ &= \chi^{\varepsilon}\left[p + \frac{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}\right]^{-1} > 0.\end{aligned}$$

Consequently, the partial derivatives $\partial p(\mu,\gamma)/\partial\mu$ and $\partial c(\mu,\gamma)/\partial\mu$ both depend on the following fraction:

$$\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}. \tag{D.5}$$

Next, we study the dependence of the fraction (D.5) on γ, as it directly determines the signs of the cross derivatives $\partial^2 p(\mu,\gamma)/\partial\mu\partial\gamma$ and $\partial^2 c(\mu,\gamma)/\partial\mu\partial\gamma$. It is given by:

$$\begin{aligned}\frac{\partial}{\partial\gamma}\left[\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}\right] &= \frac{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))\left[\frac{1}{\sigma}g(\tau)\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}\right]}{[f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))]^2}\\ &- \frac{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))\left[(1-G(\tau))\,\frac{1-\eta}{\eta\rho}\frac{\partial\theta(\mu,\gamma)}{\partial\gamma}(1+\xi)\theta^{\xi-1}(1+\theta^{\xi})^{-\frac{1+\xi}{\xi}-1}\right]}{[f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))]^2}\\ &- \frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}\left[\frac{\partial\theta(\mu,\gamma)}{\partial\gamma}q(\theta)^{1+\xi}\,(1-G(\tau)+f(\theta)G(\tau))\right]}{[f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))]^2}\\ &- \frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}\left[f(\theta)\left(\frac{1}{\sigma}g(\tau)+\frac{\partial\theta(\mu,\gamma)}{\partial\gamma}q(\theta)^{1+\xi}G(\tau)-f(\theta)\frac{1}{\sigma}g(\tau)\right)\right]}{[f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))]^2},\end{aligned}$$

and is proportional to:

$$\begin{aligned}
&\frac{\partial}{\partial\gamma}\left[\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}\right] \\
&\propto f(\theta)^2\frac{1}{\sigma}g(\tau)-f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))\,(1-G(\tau))\,\frac{\partial\theta(\mu,\gamma)}{\partial\gamma}(1+\xi)\frac{\theta^{\xi-1}}{1+\theta^{\xi}} \\
&\quad-(1-G(\tau))\,\frac{\partial\theta(\mu,\gamma)}{\partial\gamma}q(\theta)^{1+\xi}\,(1-G(\tau)+f(\theta)G(\tau)) \\
&\quad-(1-G(\tau))\,f(\theta)\frac{\partial\theta(\mu,\gamma)}{\partial\gamma}q(\theta)^{1+\xi}G(\tau),
\end{aligned}$$

where $g(\tau)\equiv\phi[(\log\tau-\gamma)/\sigma]$ and $\phi(.)$ is the standard normal probability density function. When the partial derivative of θ with respect to γ satisfies the following:

$$\frac{\partial\theta(\mu,\gamma)}{\partial\gamma}>\frac{f(\theta)^2\frac{1}{\sigma}g(\tau)}{(1-G(\tau))\,f(\theta)q(\theta)^{1+\xi}G(\tau)}=\frac{\theta(1+\theta^{\xi})\frac{1}{\sigma}g(\tau)}{(1-G(\tau))\,G(\tau)}>0,$$

it is straightforward to verify that:

$$\frac{\partial}{\partial\gamma}\left[\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}\right]<0,$$

and the values of the cross derivatives for price and consumption (or, equivalently, output) can thus be determined:

$$\begin{aligned}
\frac{\partial^2 p(\mu,\gamma)}{\partial\mu\partial\gamma}=&-\left\{\frac{\partial}{\partial\gamma}\left[\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}\right]-\frac{\partial p(\mu,\gamma)}{\partial\gamma}\frac{1}{p^2}\right\} \\
&\cdot\left[\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}+\frac{1}{p}\right]^{-2}\frac{1}{\mu}>0,
\end{aligned}$$

$$\begin{aligned}
\frac{\partial^2 c(\mu,\gamma)}{\partial\mu\partial\gamma}=&-\left\{\frac{\partial p(\mu,\gamma)}{\partial\gamma}-\frac{\partial}{\partial\gamma}\left[\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}\right]\left[\frac{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}\right]^{-2}\right\} \\
&\cdot\left[p+\frac{f(\theta)\,(1-G(\tau)+f(\theta)G(\tau))}{(1-G(\tau))\,\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}\right]^{-2}\chi^{\varepsilon}<0,
\end{aligned}$$

where $\partial p(\mu,\gamma)/\partial\gamma>0$ according to Equation (D.2). With the derivatives of p and c determined, deriving the rest of the derivatives in Proposition 6 becomes straightforward. First, concerning product market tightness, its partial and cross derivatives are given by:

$$\frac{\partial\theta(\mu,\gamma)}{\partial\mu}=\frac{1-\eta}{\eta\rho}\frac{\partial p(\mu,\gamma)}{\partial\mu}>0,\ \ \frac{\partial^2\theta(\mu,\gamma)}{\partial\mu\partial\gamma}=\frac{1-\eta}{\eta\rho}\frac{\partial^2 p(\mu,\gamma)}{\partial\mu\partial\gamma}>0.$$

Those corresponding to the wholesale price are:

$$\frac{\partial r(\mu,\gamma)}{\partial \mu} = \frac{\partial p(\mu,\gamma)}{\partial \mu} > 0, \ \frac{\partial^2 r(\mu,\gamma)}{\partial \mu \partial \gamma} = \frac{\partial^2 p(\mu,\gamma)}{\partial \mu \partial \gamma} > 0.$$

Lastly, regarding the matching cost and spare capacity, their derivatives are expressed as follows:

$$\frac{\partial}{\partial \mu}\left[G(\tau)l - c(\mu,l)\right] = -\frac{\partial c(\mu,\gamma)}{\partial \mu} < 0, \ \frac{\partial^2}{\partial \mu \partial \gamma}\left[G(\tau)l - c(\mu,l)\right] = -\frac{\partial^2 c(\mu,\gamma)}{\partial \mu \partial \gamma} > 0.$$

$$\frac{\partial}{\partial \mu}\left[l - c(\mu,l)\right] = -\frac{\partial c(\mu,\gamma)}{\partial \mu} < 0, \ \frac{\partial^2}{\partial \mu \partial \gamma}\left[l - c(\mu,l)\right] = -\frac{\partial^2 c(\mu,\gamma)}{\partial \mu \partial \gamma} > 0.$$

D.7. Theoretical Prediction on the Effectiveness of Monetary Policy When Productive Capacity Is Constrained

Next, we derive our theoretical prediction on the effectiveness of monetary policy in controlling inflation and output, depending on whether the productive capacity of the economy is constrained or not. Similar to the scenario in which the supply chain is disrupted (as in Proposition 6), contractionary monetary policy is more effective at taming inflation and reducing the sensitivity of output when the productive capacity is constrained. The only difference, as we will elaborate below, is that the state-dependent effects of monetary policy are no longer conditional.

Referring to our theoretical model in Section 3, we describe a change in the stance of monetary policy through a change in the money supply parameter μ. Whether the productive capacity is constrained or not is represented by a change in the fixed factor endowment of producers l. The comparative statics are summarized next.

Proposition 6′. *For any given threshold of reservation transportation cost $\tau > 0$ and parameter values relevant for monetary policy $\mu \in \mathbb{R}^+$ and productive capacity $l \in \mathbb{R}^+$, the responses of the endogenous variables to a change in monetary policy are described by the following partial derivatives:*

$$\frac{\partial c(\mu,l)}{\partial \mu} > 0, \ \frac{\partial p(\mu,l)}{\partial \mu} > 0, \ \frac{\partial \theta(\mu,l)}{\partial \mu} > 0, \ \frac{\partial r(\mu,l)}{\partial \mu} > 0,$$
$$\frac{\partial}{\partial \mu}\left[G(\tau)l - c(\mu,l)\right] < 0, \ \frac{\partial}{\partial \mu}\left[l - c(\mu,l)\right] < 0.$$

The cross derivatives that describe the variations in the responses of the endogenous variables

ascribed to the constrained productive capacity satisfy:

$$\frac{\partial^2 c(\mu,l)}{\partial\mu\partial l}>0,\ \frac{\partial^2 p(\mu,l)}{\partial\mu\partial l}<0,\ \frac{\partial^2\theta(\mu,l)}{\partial\mu\partial l}<0,\ \frac{\partial^2 r(\mu,l)}{\partial\mu\partial l}<0,$$
$$\frac{\partial^2}{\partial\mu\partial l}\left[G(\tau)l-c(\mu,l)\right]<0,\ \frac{\partial^2}{\partial\mu\partial l}\left[l-c(\mu,l)\right]<0,$$

where $c, p, \theta, r, G(\tau)l-c$, and $l-c$ represent consumption (or, equivalently, output), price, product market tightness, wholesale price, matching cost, and spare capacity, respectively.

Proof. Once again, we rely on the implicit function theorem, to derive the partial derivative $\partial p(\mu,l)/\partial\mu$. Consider the function $\mathbb{T}:[p^{min},+\infty)\times\mathbb{R}^+\times\mathbb{R}^+\to\mathbb{R}$:

$$\begin{aligned}\mathbb{T}(p,\mu,l)&=\chi^{\varepsilon}\frac{\mu}{G(\tau)l}-\frac{f(\theta)p}{1-G(\tau)+f(\theta)G(\tau)}\\&=\chi^{\varepsilon}\frac{\mu}{G(\tau)l}-\frac{\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^\tau G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}p}{1-G(\tau)+\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^\tau G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}G(\tau)}.\end{aligned}$$

Subsequently, assuming the existence of a tuple $(p_0,\mu_0,l_0)\in[p^{min},+\infty)\times\mathbb{R}^+\times\mathbb{R}^+$ such that $\mathbb{T}(p_0,\mu_0,l_0)=0$ and $\partial\mathbb{T}(p,\mu,l)/\partial p|_{p=p_0,\mu=\mu_0,l=l_0}\neq 0$, by the implicit function theorem, there is a neighborhood of (p_0,μ_0,l_0) such that whenever (μ,l) is sufficiently close to (μ_0,l_0), there is a unique p so that $\mathbb{T}(p,\mu,l)=0$. This assignment makes p a continuous function of μ and l. Applying implicit differentiation to $\mathbb{T}(p,\mu,l)$ around (p_0,μ_0,l_0) yields:

$$\frac{\partial p(\mu,l)}{\partial\mu}=-\frac{\partial\mathbb{T}(p,\mu,l)/\partial\mu}{\partial\mathbb{T}(p,\mu,l)/\partial p}.$$

The numerator can be written as:

$$\frac{\partial\mathbb{T}(p,\mu,l)}{\partial\mu}=\chi^{\varepsilon}\frac{1}{G(\tau)l}>0,$$

whereas the denominator is given by:

$$\frac{\partial\mathbb{T}(p,\mu,l)}{\partial p}=-\frac{(1-G(\tau))\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}p+f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}{\left(1-G(\tau)+f(\theta)G(\tau)\right)^2}<0.$$

By combining $\partial\mathbb{T}(p,\mu,l)/\partial\mu$ with $\partial\mathbb{T}(p,\mu,l)/\partial p$, we get that:

$$\frac{\partial p(\mu,l)}{\partial\mu}=\frac{1}{\mu}\left[\frac{(1-G(\tau))\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}+\frac{1}{p}\right]^{-1}>0.$$

Using the expression for the aggregate demand in Equation (25), we can also derive the partial

derivative of consumption (or, equivalently, output):

$$\frac{\partial c(\mu,\gamma)}{\partial \mu} = \chi^{\varepsilon}\left[p + \frac{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}\right]^{-1} > 0.$$

Next, we study the dependence of the partial derivatives on l by calculating the corresponding cross derivatives. For the price, it is written as:

$$\begin{aligned}\frac{\partial^2 p(\mu,l)}{\partial \mu \partial l} =& \Bigg[\frac{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}\theta^{\xi-1}(1+\xi)(1+\theta^{\xi})^{-\frac{1+\xi}{\xi}-1}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}\frac{\partial\theta(\mu,l)}{\partial l} \\ &+ \frac{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}q(\theta)^{2(1+\xi)}\left(1-G(\tau)+2f(\theta)G(\tau)\right)}{\left[f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)\right]^2}\frac{\partial\theta(\mu,l)}{\partial l} + \frac{1}{p^2}\frac{\partial p(\mu,l)}{\partial l}\Bigg] \\ &\cdot \left[\frac{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)} + \frac{1}{p}\right]^{-2}\frac{1}{\mu} < 0,\end{aligned}$$

since the partial derivatives of product market tightness and price with respect to l are both negative, i.e., $\partial\theta(\mu,l)/\partial l < 0$ and $\partial p(\mu,l)/\partial l < 0$.[11] For consumption (or, equivalently, output), it is given by:

$$\begin{aligned}\frac{\partial^2 c(\mu,l)}{\partial \mu \partial l} =& -\Bigg\{\frac{\partial p(\mu,l)}{\partial l} + \Bigg[\frac{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}\theta^{\xi-1}(1+\xi)(1+\theta^{\xi})^{-\frac{1+\xi}{\xi}-1}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}\frac{\partial\theta(\mu,l)}{\partial l} \\ &+ \frac{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}q(\theta)^{2(1+\xi)}\left(1-G(\tau)+2f(\theta)G(\tau)\right)}{\left[f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)\right]^2}\frac{\partial\theta(\mu,l)}{\partial l}\Bigg]\left[\frac{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}\right]^{-2}\Bigg\} \\ &\cdot \left[p + \frac{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}{\left(1-G(\tau)\right)\frac{1-\eta}{\eta\rho}q(\theta)^{1+\xi}}\right]^{-2}\chi^{\varepsilon} > 0.\end{aligned}$$

With the derivatives of p and c determined, deriving the rest of the derivatives in Proposition 6′ becomes straightforward. First, concerning product market tightness, its partial and cross derivatives are given by:

$$\frac{\partial\theta(\mu,l)}{\partial\mu} = \frac{1-\eta}{\eta\rho}\frac{\partial p(\mu,l)}{\partial\mu} > 0, \quad \frac{\partial^2\theta(\mu,l)}{\partial\mu\partial l} = \frac{1-\eta}{\eta\rho}\frac{\partial^2 p(\mu,l)}{\partial\mu\partial l} < 0.$$

As for the wholesale price, the derivatives are written as follows:

$$\frac{\partial r(\mu,l)}{\partial\mu} = \frac{\partial p(\mu,l)}{\partial\mu} > 0, \quad \frac{\partial^2 r(\mu,l)}{\partial\mu\partial l} = \frac{\partial^2 p(\mu,l)}{\partial\mu\partial l} < 0.$$

[11] Recall our discussion in Section 3.4 and Appendix D.4, both product market tightness and price will increase following an adverse shock to productive capacity.

Lastly, regarding the matching cost and spare capacity, their derivatives are given by:

$$\frac{\partial}{\partial \mu}\left[G(\tau)l - c(\mu,l)\right] = -\frac{\partial c(\mu,l)}{\partial \mu} < 0,\ \frac{\partial^2}{\partial \mu \partial l}\left[G(\tau)l - c(\mu,l)\right] = -\frac{\partial^2 c(\mu,l)}{\partial \mu \partial l} < 0,$$

$$\frac{\partial}{\partial \mu}\left[l - c(\mu,l)\right] = -\frac{\partial c(\mu,l)}{\partial \mu} < 0,\ \frac{\partial^2}{\partial \mu \partial l}\left[l - c(\mu,l)\right] = -\frac{\partial^2 c(\mu,l)}{\partial \mu \partial l} < 0.$$

■

D.8. Fixed Price Aggregate Supply

In contrast to the flexible price aggregate supply discussed in the main text, we consider here an alternative pricing mechanism in which the price of goods is fixed, while the reservation transportation cost can vary. As such, the equilibrium product market tightness becomes a function of the reservation transportation cost $\bar{z}$ only:

$$\theta(\bar{z}) = \frac{1-\eta}{\eta\rho}\left(\kappa - \bar{z} + \beta \int_0^{\bar{z}} G(z')dz'\right), \tag{D.6}$$

where $\kappa \in (0, +\infty)$ is an arbitrary price. Its key analytical properties are summarized below:

Proposition 1′. *For an arbitrary price $\kappa \in (0, +\infty)$, the equilibrium product market tightness $\theta(\bar{z})$ has the following properties:*

1. $\lim_{\bar{z}\to 0^+} \theta(\bar{z}) = (1-\eta)\kappa/(\eta\rho)$ *and* $\theta(\bar{z}^{max}) = 0$, *where* $\bar{z}^{max}$ *satisfies:*

$$\kappa - \bar{z}^{max} + \beta \int_0^{\bar{z}^{max}} G(z')dz' = 0;$$

2. $\theta(\bar{z})$ *is strictly decreasing on* $(0, \bar{z}^{max}]$; *and*

3. $\theta(\bar{z})$ *is convex on* $(0, \bar{z}^{max}]$.

Proof. The first property is obvious from the definition of $\bar{z}^{max}$, as θ cannot be negative. In terms of the second and third properties, we derive that:

$$\frac{d\theta(\bar{z})}{d\bar{z}} = \frac{1-\eta}{\eta\rho}\left(-1 + \beta G(\bar{z})\right) < 0,$$

$$\frac{d^2\theta(\bar{z})}{d\bar{z}^2} = \frac{(1-\eta)\beta}{\eta\rho}\frac{1}{\bar{z}\sigma}\phi\left(\frac{\log\bar{z} - \gamma}{\sigma}\right) > 0,$$

where $\phi(.)$ is the standard normal probability density function. Hence, the equilibrium product market tightness $\theta(\bar{z})$ is strictly decreasing and convex on $(0, \bar{z}^{max}]$. ■

Proposition 1′ establishes that the equilibrium product market tightness strictly decreases with the reservation transportation cost. Specifically, an increase in the reservation transportation cost raises the expected total surplus, since matches are less likely to dissolve in the subsequent period. However, this positive effect is outweighed by the decrease in the profit margin, resulting in a net negative impact on the total surplus.

Using the equilibrium product market tightness in Equation (D.6), we define the fixed price aggregate supply as follows, and its key analytical properties are summarized in Proposition 2′.

Definition 2″. *For an arbitrary price $\kappa \in (0, +\infty)$, the fixed price aggregate supply c_s^{fix} is the function of reservation transportation cost $\bar{z}$ defined by:*

$$c_s^{fix}(\bar{z}) = \frac{\left\{1 + \left[\frac{1-\eta}{\eta\rho}\left(\kappa - \bar{z} + \beta\int_0^{\bar{z}} G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}} G(\bar{z})l}{1 - G(\bar{z}) + \left\{1 + \left[\frac{1-\eta}{\eta\rho}\left(\kappa - \bar{z} + \beta\int_0^{\bar{z}} G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}} G(\bar{z})}, \tag{D.7}$$

for all $\bar{z} \in (0, \bar{z}^{max}]$, where $\bar{z}^{max}$ satisfies:

$$\kappa - \bar{z}^{max} + \beta\int_0^{\bar{z}^{max}} G(z')dz' = 0. \tag{D.8}$$

Proposition 2′. *The fixed price aggregate supply c_s^{fix} has the following properties:*

1. $\lim_{\bar{z}\to 0^+} c_s^{fix}(\bar{z}) = 0$;
2. $c_s^{fix}(\bar{z}^{max}) = 0$; *and*
3. *There exists at least one $\bar{z}^* \in (0, \bar{z}^{max}]$ such that $dc_s^{fix}(\bar{z})/d\bar{z}|_{\bar{z}=\bar{z}^*} = 0$.*

Proof. It is straightforward to prove the first property. When $\bar{z} \to 0^+$, $\lim_{\bar{z}\to 0^+}\theta(\bar{z}) = (1-\eta)\kappa/(\eta\rho)$, $\lim_{\bar{z}\to 0^+} f(\theta(\bar{z})) = \{1 + [(1-\eta)\kappa/(\eta\rho)]^{-\xi}\}^{-1/\xi} > 0$. At the same time, when $\bar{z} \to 0^+$, $\lim_{\bar{z}\to 0^+} G(\bar{z}) = 0$. Therefore, $\lim_{\bar{z}\to 0^+} c_s^{fix}(\bar{z}) = 0$. In terms of the second property, it is obvious from the definition of $\bar{z}^{max}$, together with $f(0) = 0$. Regarding the last property, the derivative of c_s^{fix} with respect to $\bar{z}$ can be expressed as follows:

$$\begin{aligned}\frac{dc_s^{fix}(\bar{z})}{d\bar{z}} =& \left(1 - G(\bar{z}) + f(\theta)G(\bar{z})\right)^{-2} \\ &\cdot \Bigg[\underbrace{\left(1 - G(\bar{z})\right)\frac{1-\eta}{\eta\rho}\left(-1 + \beta G(\bar{z})\right)\theta^{-\xi-1} f(\theta)^{1+\xi} G(\bar{z})l}_{\text{Product Market Tightness}} + \underbrace{f(\theta)\frac{1}{\bar{z}\sigma}g(\bar{z})l}_{\text{Separation Margin}}\Bigg],\end{aligned}$$

Product Market Tightness Channel < 0; Separation Margin Channel > 0

where $G(\bar{z}) \equiv \Phi[(\log \bar{z} - \gamma)/\sigma]$, $g(\bar{z}) \equiv \phi[(\log \bar{z} - \gamma)/\sigma]$, while $\Phi(.)$ and $\phi(.)$ are the standard normal cumulative density function and probability density function respectively. The product market tightness channel is negative, because a higher reservation transportation cost would reduce the total surplus to be shared between producers and retailers at the margin, hence dampening the incentives for retailers to visit producers, leading to a slack product market as well as a lower aggregate supply. On the other hand, the separation margin channel is positive, because a larger proportion of matches that would otherwise have been dismissed could now continue, hence contributing to a higher aggregate supply. The fixed price aggregate supply is determined jointly by these two channels, and the extent to which one channel dominates the other depends on both the parameter values and the reservation transportation cost itself.

When $\bar{z} \to 0^+$, we can show that the slope of the fixed price aggregate supply curve is positive:

$$\lim_{\bar{z}\to 0^+} \frac{dc_s^{fix}(\bar{z})}{d\bar{z}} = \left\{1 + \left[\frac{(1-\eta)\kappa}{\eta\rho}\right]^{-\xi}\right\}^{-\frac{1}{\xi}} \lim_{\bar{z}\to 0^+} \frac{1}{\bar{z}\sigma} g(\bar{z}) l > 0,$$

since the probability density function of a log-normal distribution is always positive. When $\bar{z} \to \bar{z}^{max}$, the slope becomes negative:

$$\lim_{\bar{z}\to \bar{z}^{max}} \frac{dc_s^{fix}(\bar{z})}{d\bar{z}} = \frac{1}{1 - G(\bar{z}^{max})} \frac{1-\eta}{\eta\rho} \left(-1 + \beta G(\bar{z}^{max})\right) G(\bar{z}^{max}) l < 0.$$

Consider $\varepsilon > 0$ such that $dc_s^{fix}(\bar{z})/d\bar{z}|_{\bar{z}=\varepsilon}$ and $\lim_{\bar{z}\to 0^+} dc_s^{fix}(\bar{z})/d\bar{z}$ have the same sign. By the intermediate value theorem, since $dc_s^{fix}(\bar{z})/d\bar{z}$ is continuous on $[\varepsilon, \bar{z}^{max}]$, there must exist at least one $\bar{z}^* \in [\varepsilon, \bar{z}^{max}]$ such that $dc_s^{fix}(\bar{z})/d\bar{z}|_{\bar{z}=\bar{z}^*} = 0$. Since $[\varepsilon, \bar{z}^{max}]$ is a sub-interval of $(0, \bar{z}^{max}]$, the last property thus holds. ■

To plot the fixed price aggregate supply, we also need to pin down its curvature. Since the value of the second derivative of c_s^{fix} cannot be determined analytically, we resort to numerical methods for an approximation. Figure D.2 plots the fixed price aggregate supply, along with the maximum volume of goods supply if the search and matching frictions were absent, and productive capacity. As observed, the non-monotonic behavior of the fixed price aggregate supply clearly illustrates the two aforementioned counteracting channels at play. Specifically, when the reservation transportation cost is relatively low, the separation margin channel dominates the product market tightness channel, and *vice versa*. Consequently, there exists a level of reservation transportation cost at which the aggregate supply is maximized. This behavior is akin to that

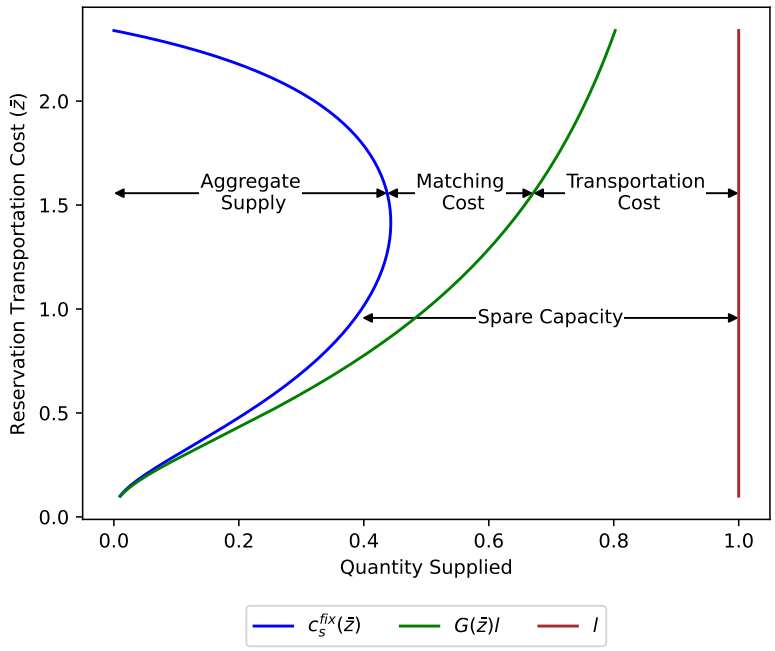

Figure D.2: Supply Side of the Economy When the Price of Goods Is Fixed

Notes. The figure plots the fixed price aggregate supply $c_s^{fix}(\bar{z})$, the maximum volume of goods supply if the search and matching frictions were absent $G(\bar{z})l$, and productive capacity l for specific values of $\bar{z}$. The difference between $c_s^{fix}(\bar{z})$ and $G(\bar{z})l$ signifies the matching cost, while the difference between $G(\bar{z})l$ and l denotes the transportation cost (measured in units of goods). The gap between $c_s^{fix}(\bar{z})$ and l represents the spare capacity of producers. For the numerical approximation, the parameter values are set as follows: $\eta = 0.5, \rho = 0.5, \kappa = 1.2, \beta = 0.99, \xi = 2, l = 1, \gamma = 0$, and $\sigma = 1$. The value of $\bar{z}$ is derived from 1,000 evenly spaced numbers over the interval $[0.1, \bar{z}^{max}]$, where $\bar{z}^{max}$ is numerically determined based on its definition in Equation (D.8).

considered in Michaillat and Saez (2015), where both the search and matching frictions in the product market and the matching cost per visit lead to the non-standard behavior of the aggregate supply curve.[12]

In terms of the other variables of interest, as the reservation transportation cost increases, the matching cost $(G(\bar{z})l - c)$ rises, the transportation cost $(l - G(\bar{z})l)$ declines, and the spare capacity represented by $l - c$ first decreases and then increases.

D.9. Supply Chain Disruptions and Reduced Matching Efficiency

Now, we extend our theoretical model by representing an adverse supply chain disturbance as a reduction in matching efficiency when producers and retailers meet to form a match. This reduction in matching efficiency decreases the probability of retailers meeting producers, thus

[12] In contrast to our fixed price aggregate supply curve, which is plotted in the $(c, \bar{z})$ plane, the aggregate supply curve in Michaillat and Saez (2015) is plotted in the (c, x) plane, where x refers to product market tightness. See Figure I of their paper for details.

imposing larger costs on retailers to form a match. As we will elaborate below, modeling a supply chain disruption shock as a reduction in matching efficiency results in the same set of comparative statics on consumption (or, equivalently, output), price, and spare capacity as an increase in the scale parameter of the log-normal distribution of transportation costs (recall Table 1). Interestingly, while modeling the disruption as increased transportation costs does not allow for the specific determination of the effects on product market tightness and wholesale price, a reduction in matching efficiency does provide this specificity. Lastly, the state-dependent effects of a contractionary monetary policy shock in controlling inflation and output are also maintained in this alternative setting. Specifically, at a steady state where the increase in product market tightness is sufficiently large during the supply chain disruption (as stated in Equation (D.11)), the disruption intensifies the decrease (or increase) in price while mitigating the decrease (or increase) in consumption (or, equivalently, output) associated with a contractionary (or expansionary) monetary policy shock.

Recall that the matching efficiency, denoted by $A \in (0, 1]$, enters the model via the matching function:

$$\mathscr{M} = A(x_U^{-\xi} + i_U^{-\xi})^{-\frac{1}{\xi}}.$$

Due to its presence, the probability for a producer to meet a retailer becomes:

$$f(\theta) = \frac{\mathscr{M}}{x_U} = A(1 + \theta^{-\xi})^{-\frac{1}{\xi}},$$

and the probability for a retailer to meet a producer changes to:

$$q(\theta) = \frac{\mathscr{M}}{i_U} = A(1 + \theta^{\xi})^{-\frac{1}{\xi}},$$

with product market tightness still defined as $\theta \equiv i_U/x_U$. The function $f(\theta)$ satisfies $f(0) = 0$, $\lim_{\theta\to+\infty} f(\theta) = A$, and $f'(\theta) = A^{-\xi}q(\theta)^{1+\xi} > 0$, whereas $q(\theta)$ satisfies $q(0) = A$, $\lim_{\theta\to+\infty} q(\theta) = 0$, and $q'(\theta) < 0$. $f(\theta)$ and $q(\theta)$ also satisfy $f(\theta)/q(\theta) = \theta$.

Since the model structure remains the same as in the baseline setup except for the matching function, we skip the model details and directly proceed to the comparative statics, focusing on the responses of the macro aggregates to adverse shocks to aggregate demand, productive capacity, and the supply chain at the flexible price steady state defined in Definition 4.

The derivation of the comparative statics for adverse demand and capacity shocks largely

follows that in Appendix D.4, except for those related to the matching cost. Recall that the matching cost is measured by the difference between the flexible price aggregate supply and the maximum amount of goods that could be supplied in the absence of search frictions. In contrast to the baseline parameterization, the matching efficiency parameter A in this alternative setting ranges from zero to one, thereby allowing the matching cost to be expressed as follows:

$$\begin{aligned} \textit{matching cost} &= \frac{AG(\tau)}{1-(1-A)G(\tau)}l - c \\ &= \left[\frac{AG(\tau)}{1-(1-A)G(\tau)} - \frac{f(\theta)G(\tau)}{1-G(\tau)+f(\theta)G(\tau)}\right] l. \end{aligned} \tag{D.9}$$

Hence, it is straightforward to verify that the matching cost will increase in the case of an adverse demand shock, while it will decrease when there is a negative productive capacity shock.

Now, we consider the comparative statics for a supply chain disturbance. We first study the effect on price by re-visiting the function $\mathbb{T} : [p^{min}, +\infty) \times (0,1] \to \mathbb{R}$:

$$\begin{aligned} \mathbb{T}(p,A) &= \chi^{\varepsilon}\frac{\mu}{G(\tau)l} - \frac{f(\theta)p}{1-G(\tau)+f(\theta)G(\tau)} \\ &= \chi^{\varepsilon}\frac{\mu}{G(\tau)l} - \frac{A\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^{\tau}G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}p}{1-G(\tau)+A\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^{\tau}G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}G(\tau)}. \end{aligned}$$

Assuming the existence of a tuple $(p_0, A_0) \in [p^{min}, +\infty) \times (0,1]$ such that $\mathbb{T}(p_0, A_0) = 0$ and $\partial\mathbb{T}(p,A)/\partial p|_{p=p_0, A=A_0} \neq 0$, by the implicit function theorem, there is a neighborhood of (p_0, A_0) such that whenever A is sufficiently close to A_0, there is a unique p so that $\mathbb{T}(p,A) = 0$. This assignment makes p a continuous function of A. Applying implicit differentiation to $\mathbb{T}(p,A)$ around (p_0, A_0) yields:

$$\frac{dp(A)}{dA} = -\frac{\partial\mathbb{T}(p,A)/\partial A}{\partial\mathbb{T}(p,A)/\partial p}.$$

In terms of $\partial\mathbb{T}(p,A)/\partial A$, it is given by:

$$\frac{\partial\mathbb{T}(p,A)}{\partial A} = -\frac{(1-G(\tau))\left(1+\theta^{-\xi}\right)^{-\frac{1}{\xi}}p}{\left(1-G(\tau)+f(\theta)G(\tau)\right)^2} < 0.$$

As for $\partial\mathbb{T}(p,A)/\partial p$, it can be written as follows:

$$\frac{\partial\mathbb{T}(p,A)}{\partial p} = -\frac{(1-G(\tau))\frac{1-\eta}{\eta\rho}A^{-\xi}q(\theta)^{1+\xi}p + f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}{\left(1-G(\tau)+f(\theta)G(\tau)\right)^2} < 0.$$

Substituting both terms into the equation for $dp(A)/dA$ yields:

$$\frac{dp(A)}{dA} = -\frac{(1-G(\tau))\left(1+\theta^{-\xi}\right)^{-\frac{1}{\xi}}p}{(1-G(\tau))\frac{1-\eta}{\eta\rho}A^{-\xi}q(\theta)^{1+\xi}p + f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)} < 0. \tag{D.10}$$

Consequently, the price p will increase following a reduction in A. Also, given that $dp(A)/dA < 0$, it is straightforward to derive the derivatives of the other endogenous variables with respect to A, except for the matching cost, since both the aggregate supply c and the maximum amount of goods that could be supplied in the absence of search frictions $AG(\tau)l/\left[1-(1-A)G(\tau)\right]$ will change following a change in A (see Equation (D.9)). Figure D.3 visualizes the results.

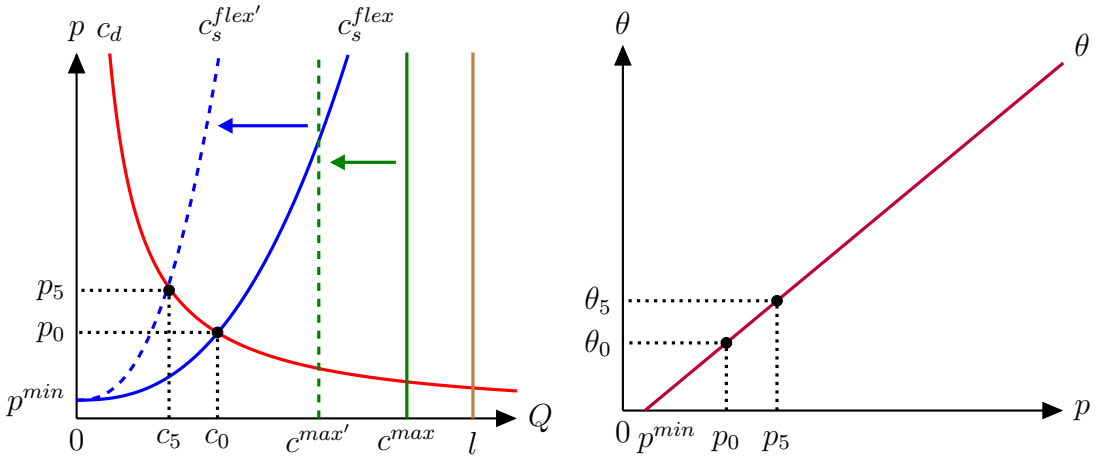

Figure D.3: An Adverse Shock to the Supply Chain, i.e., $A \downarrow$

Notes. c_s^{flex} and $c_s^{flex'}$ represent the flexible price aggregate supply before and after a reduction in matching efficiency A, respectively. c^{max} and $c^{max'}$ represent the corresponding maximum quantities of goods that could be supplied in the absence of search frictions.

In the last part, we study the state-dependent effects of a contractionary monetary policy shock when the supply chain disruption manifests as a reduction in matching efficiency. To do so, we re-define the function $\mathbb{T}$ on the joint space of $p \in [p^{min}, +\infty), \mu \in \mathbb{R}^+$, and $A \in (0,1]$:

$$\mathbb{T}(p,\mu,A) = \chi^{\varepsilon}\frac{\mu}{G(\tau)l} - \frac{A\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^\tau G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}p}{1-G(\tau)+A\left\{1+\left[\frac{1-\eta}{\eta\rho}\left(p-\tau+\beta\int_0^\tau G(z')dz'\right)\right]^{-\xi}\right\}^{-\frac{1}{\xi}}G(\tau)}.$$

Assuming the existence of a tuple $(p_0,\mu_0,A_0) \in [p^{min},+\infty)\times\mathbb{R}^+\times(0,1]$ such that $\mathbb{T}(p_0,\mu_0,A_0)=0$ and $\partial\mathbb{T}(p,\mu,A)/\partial p|_{p=p_0,\mu=\mu_0,A=A_0} \neq 0$, by the implicit function theorem, there is a neighborhood of (p_0,μ_0,A_0) such that whenever (μ,A) is sufficiently close to (μ_0,A_0), there is a unique p so that

$\mathbb{T}(p, \mu, A) = 0$. This assignment makes p a continuous function of μ and A. Applying implicit differentiation to $\mathbb{T}(p, \mu, A)$ around (p_0, μ_0, A_0) yields:

$$\frac{\partial p(\mu, A)}{\partial \mu} = -\frac{\partial \mathbb{T}(p, \mu, A)/\partial \mu}{\partial \mathbb{T}(p, \mu, A)/\partial p}.$$

The numerator is written as follows:

$$\frac{\partial \mathbb{T}(p, \mu, A)}{\partial \mu} = \chi^{\varepsilon} \frac{1}{G(\tau) l} > 0,$$

whereas the denominator is given by:

$$\frac{\partial \mathbb{T}(p, \mu, A)}{\partial p} = -\frac{(1 - G(\tau)) \frac{1-\eta}{\eta\rho} A^{-\xi} q(\theta)^{1+\xi} p + f(\theta) (1 - G(\tau) + f(\theta) G(\tau))}{(1 - G(\tau) + f(\theta) G(\tau))^2} < 0.$$

Combining terms yields the partial derivative of p with respect to μ:

$$\frac{\partial p(\mu, A)}{\partial \mu} = \frac{1}{\mu} \left[\frac{(1 - G(\tau)) \frac{1-\eta}{\eta\rho} A^{-\xi} q(\theta)^{1+\xi}}{f(\theta) (1 - G(\tau) + f(\theta) G(\tau))} + \frac{1}{p} \right]^{-1} > 0.$$

For the partial derivative of consumption (or, equivalently, output) with respect to μ, using the expression for the aggregate demand in Equation (25), we can derive that:

$$\frac{\partial c(\mu, A)}{\partial \mu} = \chi^{\varepsilon} \left[p + \frac{f(\theta) (1 - G(\tau) + f(\theta) G(\tau))}{(1 - G(\tau)) \frac{1-\eta}{\eta\rho} A^{-\xi} q(\theta)^{1+\xi}} \right]^{-1} > 0.$$

Consequently, both the cross derivatives $\partial^2 p(\mu, A)/\partial\mu\partial A$ and $\partial^2 c(\mu, A)/\partial\mu\partial A$ depend on how the following fraction,

$$\frac{(1 - G(\tau)) \frac{1-\eta}{\eta\rho} A^{-\xi} q(\theta)^{1+\xi}}{f(\theta) (1 - G(\tau) + f(\theta) G(\tau))},$$

responds to a decrease in A. Specifically, we show that the partial derivative of the fraction with respect to A is proportional to the following expression:

$$\begin{aligned}
&\frac{\partial}{\partial A} \left[\frac{(1 - G(\tau)) \frac{1-\eta}{\eta\rho} A^{-\xi} q(\theta)^{1+\xi}}{f(\theta) (1 - G(\tau) + f(\theta) G(\tau))} \right] \\
&\propto -(1 + \theta^{-\xi})^{-\frac{1}{\xi}} \left(1 - G(\tau) + A(1 + \theta^{-\xi})^{-\frac{1}{\xi}} G(\tau) \right) \theta^{\xi - 1} (1 + \xi)(1 + \theta^{\xi})^{-\frac{1+\xi}{\xi} - 1} \frac{\partial \theta(\mu, A)}{\partial A} \\
&\quad - (1 + \theta^{\xi})^{-\frac{1+\xi}{\xi}} \left(1 - G(\tau) + A(1 + \theta^{-\xi})^{-\frac{1}{\xi}} G(\tau) \right) \theta^{-(1+\xi)} (1 + \theta^{-\xi})^{-\frac{1}{\xi} - 1} \frac{\partial \theta(\mu, A)}{\partial A} \\
&\quad - (1 + \theta^{\xi})^{-\frac{1+\xi}{\xi}} (1 + \theta^{-\xi})^{-\frac{2}{\xi}} G(\tau) \\
&\quad - (1 + \theta^{\xi})^{-\frac{1+\xi}{\xi}} A(1 + \theta^{-\xi})^{-\frac{1}{\xi}} G(\tau) \theta^{-(1+\xi)} (1 + \theta^{-\xi})^{-\frac{1}{\xi} - 1} \frac{\partial \theta(\mu, A)}{\partial A}.
\end{aligned}$$

Now, suppose that the partial derivative $\partial\theta(\mu, A)/\partial A$ satisfies the condition defined below:

$$\frac{\partial\theta(\mu, A)}{\partial A} < -\frac{1+\theta^{\xi}}{A\theta^{\xi-1}}, \tag{D.11}$$

we can easily verify that the following inequality holds:

$$-(1+\theta^{-\xi})^{-\frac{1}{\xi}}A(1+\theta^{-\xi})^{-\frac{1}{\xi}}G(\tau)\theta^{\xi-1}(1+\theta^{\xi})^{-\frac{1+\xi}{\xi}-1}\frac{\partial\theta(\mu, A)}{\partial A} - (1+\theta^{\xi})^{-\frac{1+\xi}{\xi}}(1+\theta^{-\xi})^{-\frac{2}{\xi}}G(\tau) > 0,$$

and subsequently,

$$\frac{\partial}{\partial A}\left[\frac{(1-G(\tau))^{\frac{1-\eta}{\eta\rho}}A^{-\xi}q(\theta)^{1+\xi}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}\right] > 0.$$

Thus, the cross derivatives of p and c with respect to μ and A can be pinned down:

$$\begin{aligned}\frac{\partial^2 p(\mu, A)}{\partial\mu\partial A} = &-\left\{\frac{\partial}{\partial A}\left[\frac{(1-G(\tau))^{\frac{1-\eta}{\eta\rho}}A^{-\xi}q(\theta)^{1+\xi}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}\right] - \frac{\partial p(\mu, A)}{\partial A}\frac{1}{p^2}\right\} \\ &\cdot\left[\frac{(1-G(\tau))^{\frac{1-\eta}{\eta\rho}}A^{-\xi}q(\theta)^{1+\xi}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)} + \frac{1}{p}\right]^{-2}\frac{1}{\mu} < 0,\end{aligned}$$

$$\begin{aligned}\frac{\partial^2 c(\mu, A)}{\partial\mu\partial A} = &-\left\{\frac{\partial p(\mu, A)}{\partial A} - \frac{\partial}{\partial A}\left[\frac{(1-G(\tau))^{\frac{1-\eta}{\eta\rho}}A^{-\xi}q(\theta)^{1+\xi}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}\right]\left[\frac{(1-G(\tau))^{\frac{1-\eta}{\eta\rho}}A^{-\xi}q(\theta)^{1+\xi}}{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}\right]^{-2}\right\} \\ &\cdot\left[p + \frac{f(\theta)\left(1-G(\tau)+f(\theta)G(\tau)\right)}{(1-G(\tau))^{\frac{1-\eta}{\eta\rho}}A^{-\xi}q(\theta)^{1+\xi}}\right]^{-2}\chi^{\varepsilon} > 0,\end{aligned}$$

where $\partial p(\mu, A)/\partial A < 0$ according to Equation (D.10). Coupled with the partial derivatives $\partial p(\mu, A)/\partial\mu > 0$ and $\partial c(\mu, A)/\partial\mu > 0$, they jointly imply that the supply chain disruption, represented by a reduction in matching efficiency, intensifies the decrease (or increase) in price while dampening the decrease (or increase) in consumption (or, equivalently, output) associated with a contractionary (or expansionary) monetary policy shock.

In terms of the other variables of interest, their behaviors can also be determined. The partial and cross derivatives for product market tightness are given by:

$$\frac{\partial\theta(\mu, A)}{\partial\mu} = \frac{1-\eta}{\eta\rho}\frac{\partial p(\mu, A)}{\partial\mu} > 0, \quad \frac{\partial^2\theta(\mu, A)}{\partial\mu\partial A} = \frac{1-\eta}{\eta\rho}\frac{\partial^2 p(\mu, A)}{\partial\mu\partial A} < 0.$$

For the wholesale price, its derivatives are expressed as follows:

$$\frac{\partial r(\mu, A)}{\partial\mu} = \frac{\partial p(\mu, A)}{\partial\mu} > 0, \quad \frac{\partial^2 r(\mu, A)}{\partial\mu\partial A} = \frac{\partial^2 p(\mu, A)}{\partial\mu\partial A} < 0.$$

Lastly, in terms of the matching cost, its partial and cross derivatives are given by:

$$\frac{\partial}{\partial \mu}\left[\frac{AG(\tau)}{1-(1-A)G(\tau)}l - c(\mu, A)\right] = -\frac{\partial c(\mu, A)}{\partial \mu} < 0,$$
$$\frac{\partial^2}{\partial \mu \partial A}\left[\frac{AG(\tau)}{1-(1-A)G(\tau)}l - c(\mu, A)\right] = -\frac{\partial^2 c(\mu, A)}{\partial \mu \partial A} < 0,$$

and those corresponding to the spare capacity are:

$$\frac{\partial}{\partial \mu}\left[l - c(\mu, A)\right] = -\frac{\partial c(\mu, A)}{\partial \mu} < 0,$$
$$\frac{\partial^2}{\partial \mu \partial A}\left[l - c(\mu, A)\right] = -\frac{\partial^2 c(\mu, A)}{\partial \mu \partial A} < 0.$$

In summary, these derivatives suggest that during a supply chain disruption, the responses of product market tightness and wholesale price become more pronounced. In contrast, the responses of matching cost and spare capacity are less pronounced.

D.10. Convergence Dynamics

In this appendix, we explore the convergence dynamics of our model from the initial steady state to the new steady state after an aggregate, unexpected, and permanent shock to aggregate demand, productive capacity, and the supply chain.

The dynamics of our model involve the evolution of two key variables – the number of matched producers ($x_{M,t}$) and product market tightness (θ_t) – following each exogenous shock. Two transition equations govern these variables:

$$x_{M,t+1} = A(1+\theta_t^{-\xi})^{-\frac{1}{\xi}}\Phi\left(\frac{\log\tau-\gamma}{\sigma}\right) + \left[1 - A(1+\theta_t^{-\xi})^{-\frac{1}{\xi}}\right]\Phi\left(\frac{\log\tau-\gamma}{\sigma}\right)x_{M,t}, \tag{D.12}$$

and

$$\theta_t = \frac{1-\eta}{\eta\rho}\left[\frac{\chi^{\varepsilon}\mu}{x_{M,t}l} - \tau + \beta\int_0^{\tau}\Phi\left(\frac{\log z'-\gamma}{\sigma}\right)dz'\right], \tag{D.13}$$

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

To examine the convergence dynamics, we first calculate the initial and new steady states of the model before and after the change in the respective parameter. This process involves setting $x_{M,t+1} = x_{M,t}$ and simultaneously solving Equations (D.12) and (D.13). Subsequently,

assuming that the system starts at the initial steady state, we analyze the rate at which the system converges to the new steady state following each of the shocks of interest. Additionally, we investigate whether this convergence is monotonic.

Table D.1: Baseline Calibration

Parameter	Description	Value
η	Bargaining power of producers	0.5
ρ	Fixed search cost	0.5
χ	Taste for consumption	1
ε	Elasticity of substitution between c and m/p	2
μ	Nominal money supply	10
l	Productive capacity	1
τ	Reservation transportation cost	7.25
β	Discount factor	0.99
γ	Scale parameter of $G(\cdot)$	1
σ	Shape parameter of $G(\cdot)$	1
A	Matching efficiency	1
ξ	Elasticity of substitution between x_U and i_U	2

We calibrate the model at a bi-weekly frequency using U.S. data, meaning that two periods in our theoretical model correspond to one lag in our monthly SVAR. The baseline parameter values, which are standard in the literature, are summarized in Table D.1. In particular, we follow Fernández-Villaverde et al. (2024) by setting the bargaining power of the producer, η, to 0.5 to evenly split the total surplus from matching. In addition, we set $\tau = 7.25$ to yield a steady-state value of spare capacity at 16.5%, reflecting the average spare capacity rate of the top five exporting countries to the U.S. (namely, Mexico, Canada, China, Germany, and Japan), weighted by the U.S. imports of goods from each country, during the pre-pandemic period of 2017-2019. Specifically, we estimate such an import-weighted average spare capacity rate using the following formula:

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

where $\mathscr{C} \equiv \{Mexico, Canada, China, Germany, Japan\}$, $Import_{i,t}$ denotes the U.S. imports of goods by customs basis from country i in month t, and $CapacityUtilization_{i,t}$ represents the capacity utilization rate for country i in month t. The monthly time series for U.S. imports

from each of the five exporting countries can be retrieved directly from the FRED database. The capacity utilization rates for Mexico and Japan are readily available from official sources. In contrast, for Canada, China, and Germany, these rates are derived by interpolating the corresponding quarterly series using the Chow-Lin method (Chow and Lin, 1971), based on each country's industrial production index. The underlying data for the capacity utilization rates and industrial production indices are either retrieved directly or computed using available data from the respective government authorities or statistics agencies: the National Institute of Statistics and Geography for Mexico, Statistics Canada for Canada, the National Bureau of Statistics for China, the Ifo Institute for Economic Research for Germany, and the Ministry of Economy, Trade and Industry for Japan. Lastly, the number of matched producers in the first iteration ($x_{M,1}$) is set to match its initial steady-state value.

D.10.1. An Adverse Shock to Aggregate Demand

We first consider an adverse shock to aggregate demand, represented either by a 10% decrease in the money supply μ from 10 to 9, or by a 10% decrease in the preference for consumption χ from 1 to 0.9. Tables D.2 and D.3 summarize the convergence dynamics of all variables of interest for these two cases, respectively. Meanwhile, Figures D.4 and D.5 plot the dynamics specific to consumption (or, equivalently, output) and price. In both scenarios, the system's convergence to the new steady state is almost instantaneous, requiring only one iteration (the model does not have a variable like capital that could generate persistence). Also, the convergence dynamics are monotonic.

Table D.2: Convergence Dynamics: An Adverse Shock to Aggregate Demand ($\mu = 10 \to 9$)

Iteration	# Matched Producers x_M	Product Market Tightness θ	Consumption (or Output) c	Price p	Wholesale Price r	Matching Cost $\frac{AG(\tau)}{1-(1-A)G(\tau)}l - c$	Spare Capacity $l - c$
Initial Steady State	0.836469	17.051289	0.836469	11.955011	13.865328	0.000235	0.163531
1	0.836469	14.660286	0.836469	10.759510	12.669827	0.000235	0.163531
New Steady State	0.836387	14.662410	0.836387	10.760572	12.670888	0.000317	0.163613

Notes. The values of the endogenous variables at the two steady states are calculated by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (D.12) and (D.13) simultaneously. In this process, the parameter μ is set to 10 in the initial steady state and adjusted to 9 in the new steady state (i.e., a 10% decrease). For the first iteration, $x_{M,1}$ is set equal to its value in the initial steady state, and μ is decreased to 9.

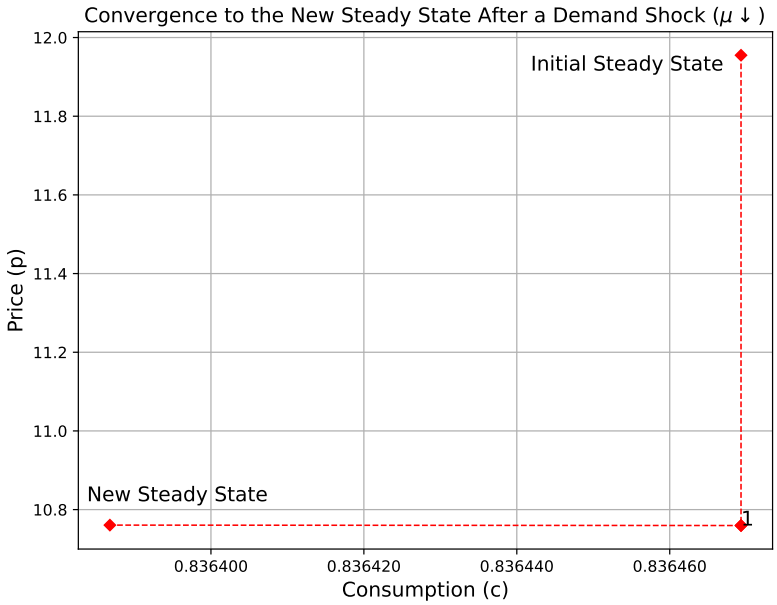

Figure D.4: Convergence Dynamics of Consumption and Price: An Adverse Shock to Aggregate Demand ($\mu = 10 \to 9$)

Notes. The figure plots the convergence dynamics of consumption (or, equivalently, output) and price from their initial steady-state values to their new steady-state values, following an adverse shock to aggregate demand (where μ decreases by 10% from 10 to 9). The iteration numbers are marked on the corresponding dots.

Table D.3: Convergence Dynamics: An Adverse Shock to Aggregate Demand ($\chi = 1 \to 0.9$)

Iteration	# Matched Producers x_M	Product Market Tightness θ	Consumption (or Output) c	Price p	Wholesale Price r	Matching Cost $\frac{AG(\tau)}{1-(1-A)G(\tau)}l - c$	Spare Capacity $l - c$
Initial Steady State	0.836469	17.051289	0.836469	11.955011	13.865328	0.000235	0.163531
1	0.836469	12.508384	0.836469	9.683559	11.593876	0.000235	0.163531
New Steady State	0.836269	12.513028	0.836269	9.685881	11.596197	0.000435	0.163731

Notes. The values of the endogenous variables at the two steady states are calculated by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (D.12) and (D.13) simultaneously. In this process, the parameter χ is set to 1 in the initial steady state and adjusted to 0.9 in the new steady state (i.e., a 10% decrease). For the first iteration, $x_{M,1}$ is set equal to its value in the initial steady state, and χ is decreased to 0.9.

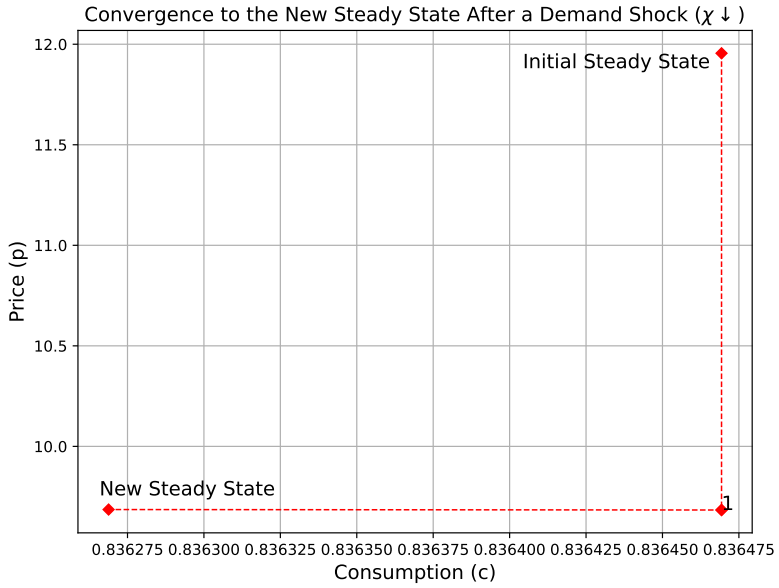

Figure D.5: Convergence Dynamics of Consumption and Price: An Adverse Shock to Aggregate Demand ($\chi = 1 \rightarrow 0.9$)

Notes. The figure plots the convergence dynamics of consumption (or, equivalently, output) and price from their initial steady-state values to their new steady-state values, following an adverse shock to aggregate demand (where χ decreases by 10% from 1 to 0.9). The iteration numbers are marked on the corresponding dots.

D.10.2. An Adverse Shock to Productive Capacity

Next, we consider an adverse shock to productive capacity, represented by a 10% decrease in the fixed factor endowment of producers l from 1 to 0.9. The convergence dynamics, as observed in Table D.4 and Figure D.6, occur almost instantaneously, requiring only one iteration to reach the new steady state. Furthermore, the process is monotonic.

Table D.4: Convergence Dynamics: An Adverse Shock to Productive Capacity ($l = 1 \to 0.9$)

Iteration	# Matched Producers x_M	Product Market Tightness θ	Consumption (or Output) c	Price p	Wholesale Price r	Matching Cost $\frac{AG(\tau)}{1-(1-A)G(\tau)}l - c$	Spare Capacity $l - c$
Initial Steady State	0.836469	17.051289	0.836469	11.955011	13.865328	0.000235	0.163531
1	0.836469	19.707958	0.752822	13.283346	15.193662	0.000211	0.147178
New Steady State	0.836528	19.706088	0.752875	13.282411	15.192727	0.000158	0.147125

Notes. The values of the endogenous variables at the two steady states are calculated by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (D.12) and (D.13) simultaneously. In this process, the parameter l is set to 1 in the initial steady state and adjusted to 0.9 in the new steady state (i.e., a 10% decrease). For the first iteration, $x_{M,1}$ is set equal to its value in the initial steady state, and l is decreased to 0.9.

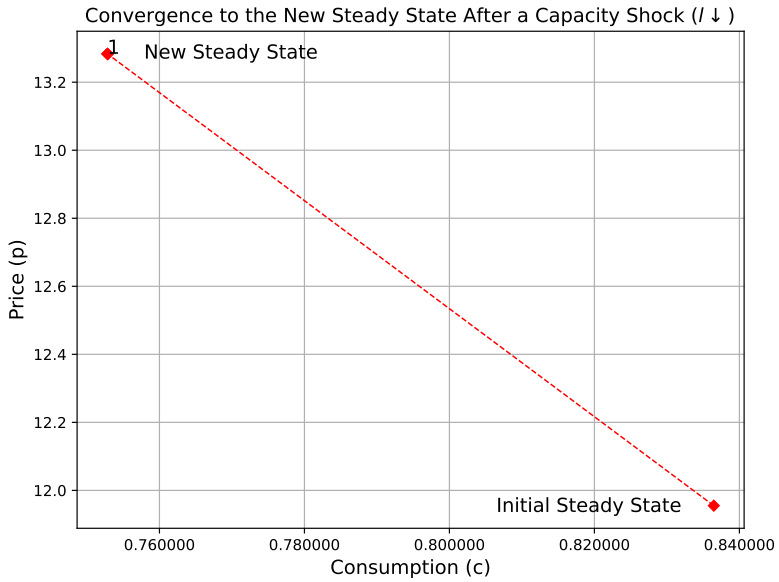

Figure D.6: Convergence Dynamics of Consumption and Price: An Adverse Shock to Productive Capacity ($l = 1 \to 0.9$)

Notes. The figure plots the convergence dynamics of consumption (or, equivalently, output) and price from their initial steady-state values to their new steady-state values, following an adverse shock to productive capacity (where l decreases by 10% from 1 to 0.9). The iteration numbers are marked on the corresponding dots.

D.10.3. An Adverse Shock to the Supply Chain

Lastly, we consider an adverse shock to the supply chain, represented either by a 10% increase in the scale parameter of the log-normal distribution of transportation costs, γ, from 1 to 1.1 (as illustrated in Table D.5 and Figure D.7), or by a 10% decrease in the matching efficiency A from 1 to 0.9 (as detailed in Table D.6 and Figure D.8). The convergence to the new steady state is achieved quickly, with the system requiring two iterations to reach the new steady state in the case of an increase in γ, and four iterations in the case of a decrease in A. Furthermore, as with the cases of adverse shocks to aggregate demand and productive capacity, the convergence dynamics are monotonic.

It is also noteworthy in Table D.5 that, in the scenario of a 10% increase in γ, the resulting price increase is sufficiently large to raise product market tightness in the new steady state. However, in an alternative scenario with only a 0.0001% increase in γ, as shown in Table D.7, the resulting price increase does not lead to greater tightness in the product market. This observation corroborates our discussion in Section 3.4, which suggests that adverse disturbances to the supply chain can either tighten or loosen product market tightness, depending on the relative magnitudes of the price increase and the decrease in the expected total surplus due to higher transportation costs.

Table D.5: Convergence Dynamics: An Adverse Shock to the Supply Chain ($\gamma = 1 \to 1.1$)

Iteration	# Matched Producers x_M	Product Market Tightness θ	Consumption (or Output) c	Price p	Wholesale Price r	Matching Cost $\frac{AG(\tau)}{1-(1-A)G(\tau)}l - c$	Spare Capacity $l - c$
Initial Steady State	0.836469	17.051289	0.836469	11.955011	13.865328	0.000235	0.163531
1	0.836469	16.610626	0.836469	11.955011	13.755162	-0.025628	0.163531
2	0.810602	17.373630	0.810602	12.336513	14.136664	0.000240	0.189398
New Steady State	0.810588	17.374059	0.810588	12.336728	14.136879	0.000254	0.189412

Notes. The values of the endogenous variables at the two steady states are calculated by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (D.12) and (D.13) simultaneously. In this process, the parameter γ is set to 1 in the initial steady state and adjusted to 1.1 in the new steady state (i.e., a 10% increase). For the first iteration, $x_{M,1}$ is set equal to its value in the initial steady state, and γ is increased to 1.1.

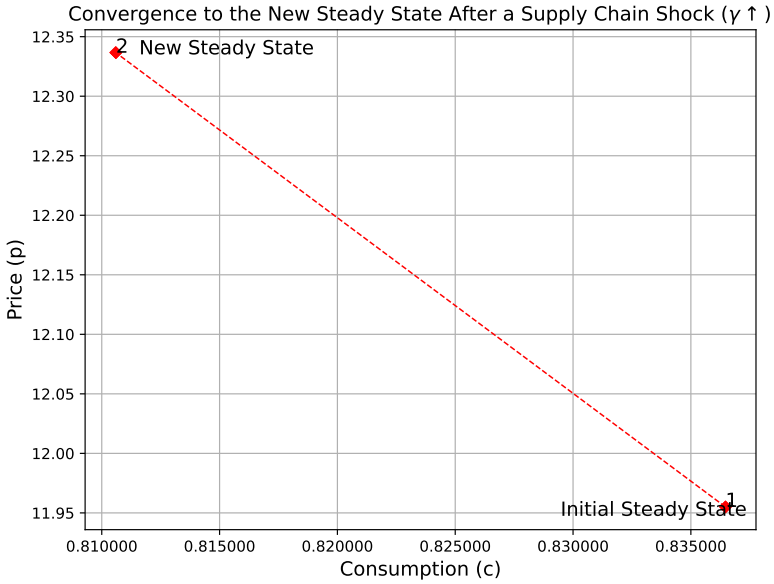

Figure D.7: Convergence Dynamics of Consumption and Price: An Adverse Shock to the Supply Chain ($\gamma = 1 \to 1.1$)

Notes. The figure plots the convergence dynamics of consumption (or, equivalently, output) and price from their initial steady-state values to their new steady-state values, following an adverse shock to the supply chain (where γ increases by 10% from 1 to 1.1). The iteration numbers are marked on the corresponding dots.

Table D.6: Convergence Dynamics: An Adverse Shock to the Supply Chain ($A = 1 \to 0.9$)

Iteration	# Matched Producers x_M	Product Market Tightness θ	Consumption (or Output) c	Price p	Wholesale Price r	Matching Cost $\frac{AG(\tau)}{1-(1-A)G(\tau)}l - c$	Spare Capacity $l - c$
Initial Steady State	0.836469	17.051289	0.836469	11.955011	13.865328	0.000235	0.163531
1	0.836469	17.051289	0.836469	11.955011	13.865328	-0.014676	0.163531
2	0.822810	17.448212	0.822810	12.153473	14.063789	-0.001017	0.177190
3	0.821660	17.482240	0.821660	12.170487	14.080803	0.000134	0.178340
4	0.821563	17.485108	0.821563	12.171921	14.082237	0.000230	0.178437
New Steady State	0.821555	17.485349	0.821555	12.172041	14.082358	0.000238	0.178445

Notes. The values of the endogenous variables at the two steady states are calculated by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (D.12) and (D.13) simultaneously. In this process, the parameter A is set to 1 in the initial steady state and adjusted to 0.9 in the new steady state (i.e., a 10% decrease). For the first iteration, $x_{M,1}$ is set equal to its value in the initial steady state, and A is decreased to 0.9.

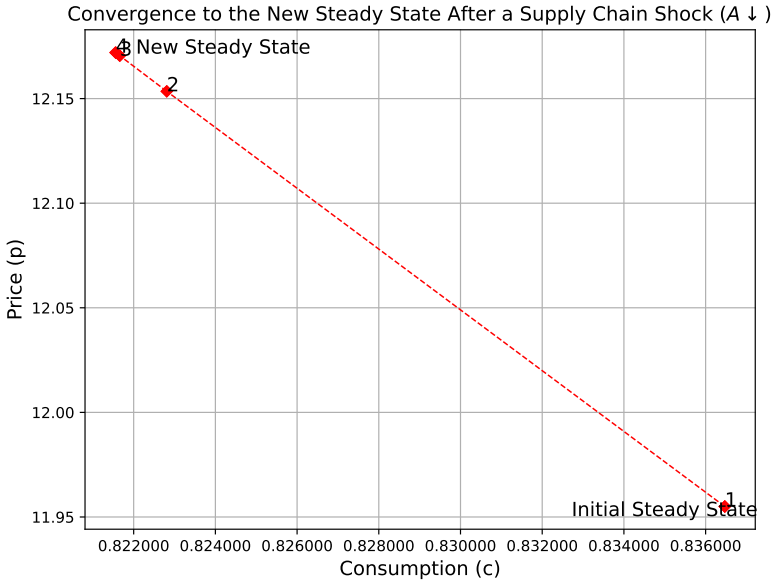

Figure D.8: Convergence Dynamics of Consumption and Price: An Adverse Shock to the Supply Chain ($A = 1 \to 0.9$)

Notes. The figure plots the convergence dynamics of consumption (or, equivalently, output) and price from their initial steady-state values to their new steady-state values, following an adverse shock to the supply chain (where A decreases by 10% from 1 to 0.9). The iteration numbers are marked on the corresponding dots.

Table D.7: Convergence Dynamics: An Adverse Shock to the Supply Chain ($\gamma = 1 \to 1.000001$)

Iteration	# Matched Producers x_M	Product Market Tightness θ	Consumption (or Output) c	Price p	Wholesale Price r	Matching Cost $\frac{AG(\tau)}{1-(1-A)G(\tau)}l - c$	Spare Capacity $l - c$
Initial Steady State	0.836469	17.051289	0.836469	11.955011	13.865328	0.000235	0.163531
New Steady State	0.836469	17.051284	0.836469	11.955011	13.865327	0.000234	0.163531

Notes. The values of the endogenous variables at the two steady states are calculated by setting $x_{M,t+1} = x_{M,t}$ and solving Equations (D.12) and (D.13) simultaneously. In this process, the parameter γ is set to 1 in the initial steady state and adjusted to 1.000001 in the new steady state (i.e., a 0.0001% increase). For the first iteration, $x_{M,1}$ is set equal to its value in the initial steady state, and γ is increased to 1.000001.

E. Robustness of SVAR Results

E.1. Dropping the Zero Restrictions

In this appendix, we conduct a robustness check of our baseline results by relaxing the zero restrictions imposed on the contemporaneous responses of the ACR index to aggregate demand and productive capacity shocks in the SVAR estimation. In the baseline estimation, we impose these zero restrictions based on our domain knowledge of the short-term rigidity of the container shipping industry (as discussed in Section 2.6). However, removing these restrictions allows us to compare our ACR index with other indices in the existing literature that track the state of the global supply chain and highlights the important role of this industry feature in achieving a tighter identification outcome. As we will elaborate later in Appendix F, other popular indices often capture factors irrelevant to supply chain disruptions – such as fuel costs, seasonal demand fluctuations, market speculation affecting shipping prices – and large, time-varying measurement errors. For example, an increase in "delivery times" – the average duration that suppliers take to provide inputs to their customers' factories – reported in the PMI could stem from either a supply chain disruption or a reduction in productive capacity.

Specifically, without the zero restrictions, we re-write the identification restrictions of the three shocks as follows:

Restriction 1′. ***An adverse shock to aggregate demand*** *leads to a negative response of real GDP, PCE goods price, product market tightness, and import price, as well as to a positive response of spare capacity at* $k = 1$.

Restriction 2′. ***An adverse shock to productive capacity*** *leads to a negative response of real GDP and spare capacity, as well as to a positive response of PCE goods price, product market tightness, and import price at* $k = 1$.

Restriction 3′. ***An adverse shock to the supply chain*** *leads to a negative response of real GDP, as well as to a positive response of PCE goods price, spare capacity, and the ACR at* $k = 1$.

The estimation is implemented with the Bayesian approach as in Arias et al. (2018, 2019, 2023). The specifications of the system, except for the identification restrictions, remain the same as those used in the baseline estimation. As seen in Figures E.1 through E.3, the IRFs remain

largely unchanged from those in the baseline estimation. Despite imposing fewer restrictions, the impulse responses of the macro aggregates still evolve as predicted by our theory. However, the confidence bands for real GDP, prices, and spare capacity are noticeably wider, especially in the cases of demand and productive capacity shocks.

Additionally, as shown in Figure E.4, productive capacity shocks replace supply chain disturbances as the primary source of unexpected fluctuations in PCE goods and import prices over the longer term. Meanwhile, all three shocks contribute to the unexpected variations in the ACR index across all horizons. Lastly, comparing the historical decomposition of U.S. goods inflation in Figure E.5 to that in Figure 14, there is little difference. The supply chain disruptions were still identified as the primary driving force behind the increase in inflation during 2021.

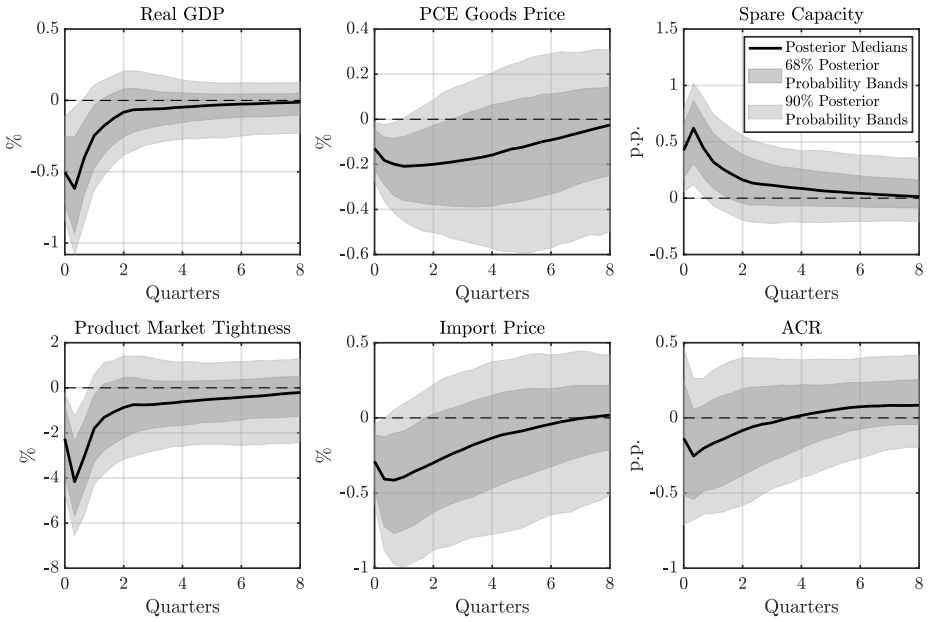

Figure E.1: IRFs to an Adverse Shock to Aggregate Demand: Dropping Zero Restrictions

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using the ACR index and Restrictions 1′, 2′, and 3′. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

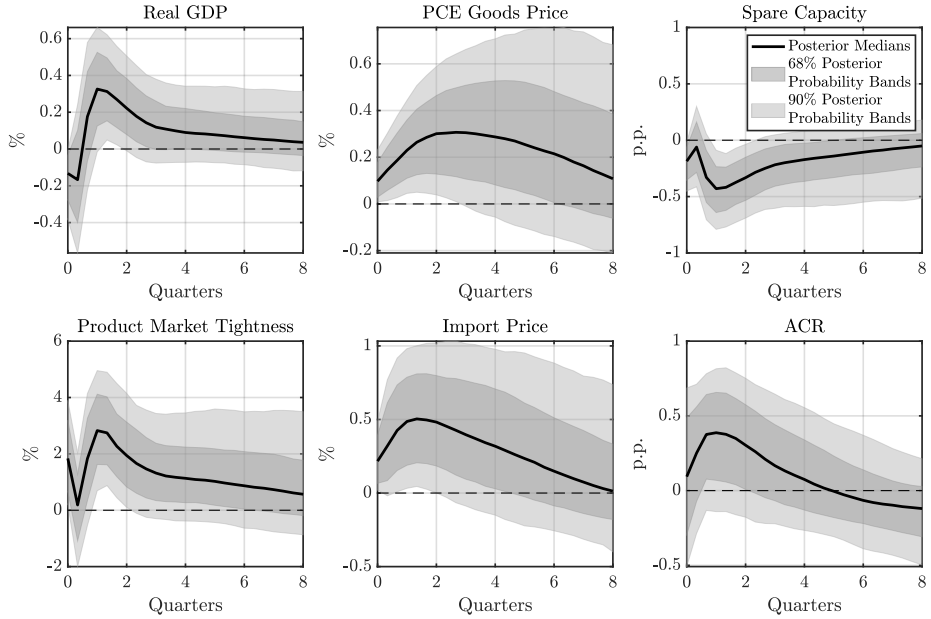

Figure E.2: IRFs to an Adverse Shock to Productive Capacity: Dropping Zero Restrictions

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using the ACR index and Restrictions 1′, 2′, and 3′. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

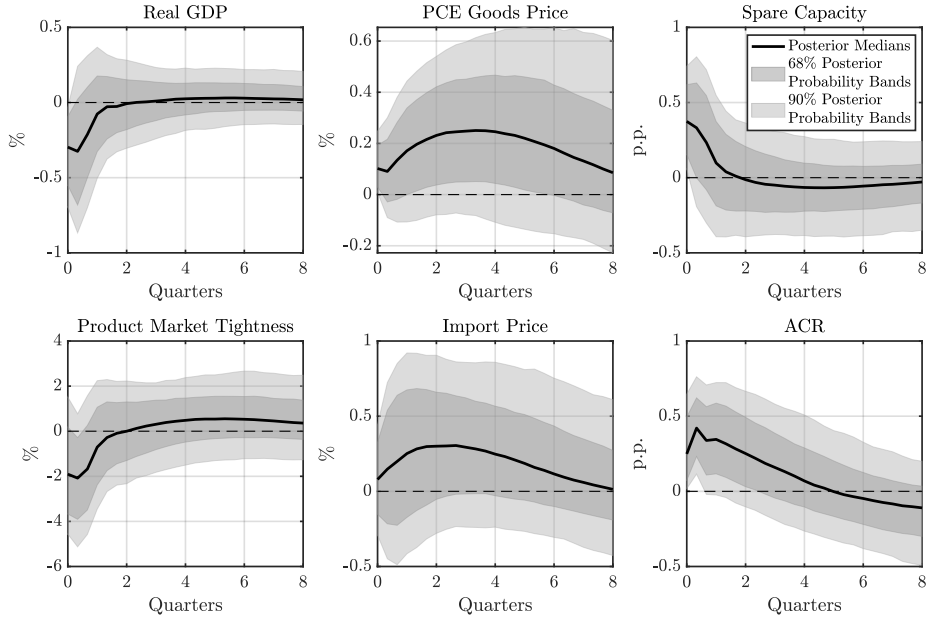

Figure E.3: IRFs to an Adverse Shock to the Supply Chain: Dropping Zero Restrictions

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using the ACR index and Restrictions 1′, 2′, and 3′. The solid line shows the point-wise posterior medians, and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

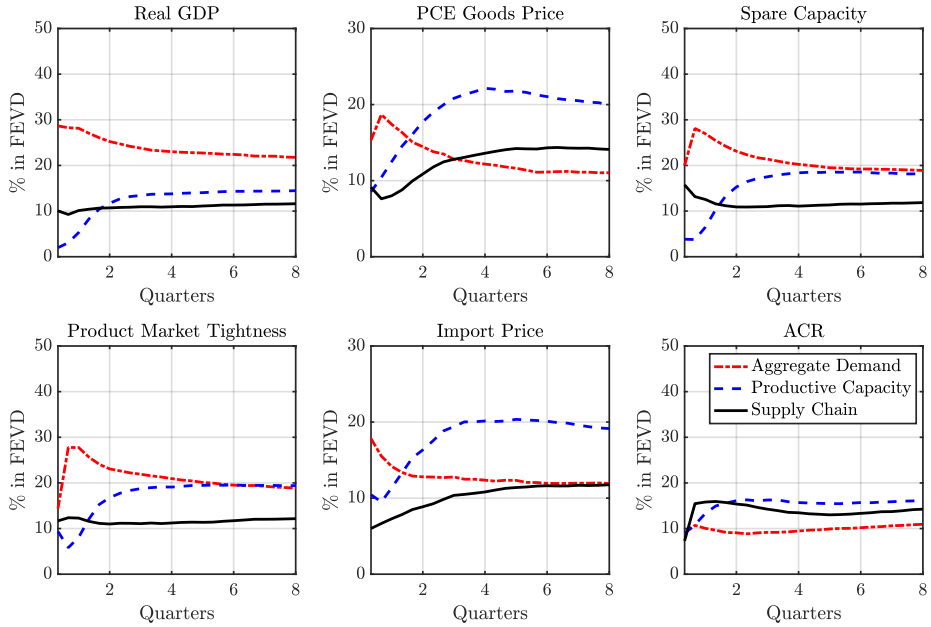

Figure E.4: FEVD from the SVAR: Dropping Zero Restrictions

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated using the ACR index and Restrictions 1′, 2′, and 3′, and based on 100,000 independent importance sampling draws.

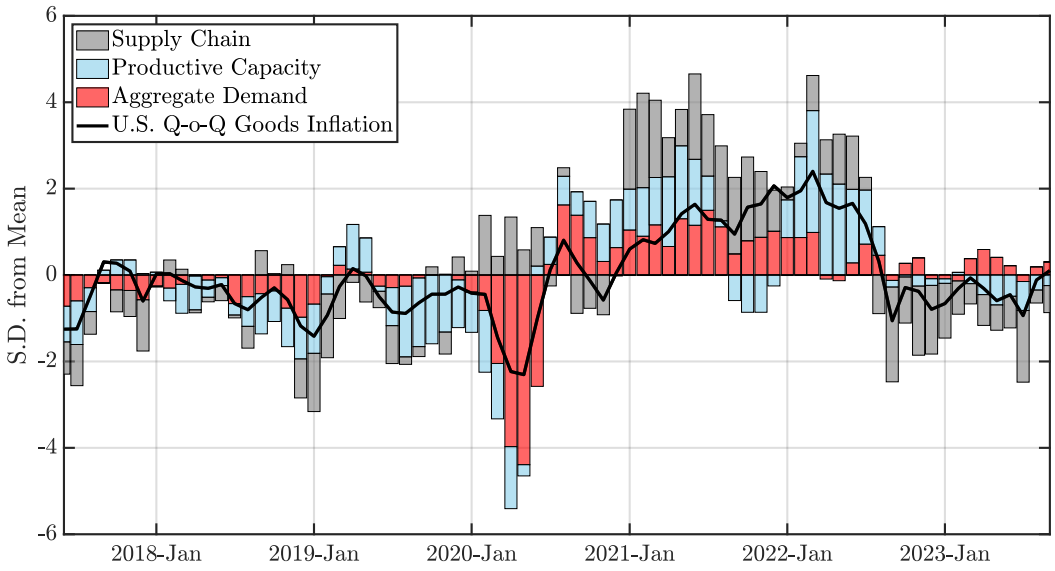

Figure E.5: HD of U.S. Quarter-on-Quarter Goods Inflation: Dropping Zero Restrictions

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with the ACR index included as the measure indicating the state of the global supply chain, and with Restrictions 1′, 2′, and 3′ imposed on the contemporaneous IRFs. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

E.2. Dropping the ACR Index

In this appendix, we check whether similar identification results for the supply chain disturbances are achieved by omitting the ACR index from the SVAR estimation. This exercise is intended to highlight the importance of including an endogenous variable that captures the state of the global supply chain when estimating the causal effects of global supply chain disruptions on macro aggregates.

Specifically, we include the same endogenous variables in the SVAR model, except for the ACR index, and impose the following identification restrictions on the contemporaneous responses of the remaining endogenous variables:

Restriction 1″. ***An adverse shock to aggregate demand*** *leads to a negative response of real GDP, PCE goods price, product market tightness, and import price, as well as to a positive response of spare capacity at* $k = 1$.

Restriction 2″. ***An adverse shock to productive capacity*** *leads to a negative response of real GDP and spare capacity, as well as to a positive response of PCE goods price, product market tightness, and import price at* $k = 1$.

Restriction 3″. ***An adverse shock to the supply chain*** *leads to a negative response of real GDP, as well as to a positive response of PCE goods price and spare capacity at* $k = 1$.

As in the baseline estimation, we use the Bayesian approach (Arias et al., 2018, 2019, 2023) to identify the causal effects of the three structural shocks. Compared to the baseline estimation, the main difference lies in the estimated effects of a supply chain disruption shock. As shown in Figure E.8, although the stagflationary effect of a supply chain disturbance is still observed, the responses of PCE goods and import prices are less precisely estimated, with the corresponding 68% confidence bands including the zero-response line. Furthermore, the FEVD in Figure E.9 indicates that supply chain disturbances account for a minimal share of unexpected fluctuations in PCE goods and import prices across all horizons, thereby underestimating the significance of supply chain disruptions in driving inflation. The historical decomposition of U.S. goods inflation also supports this point since supply chain shocks are not identified as the primary driver of inflation until the second half of 2021.

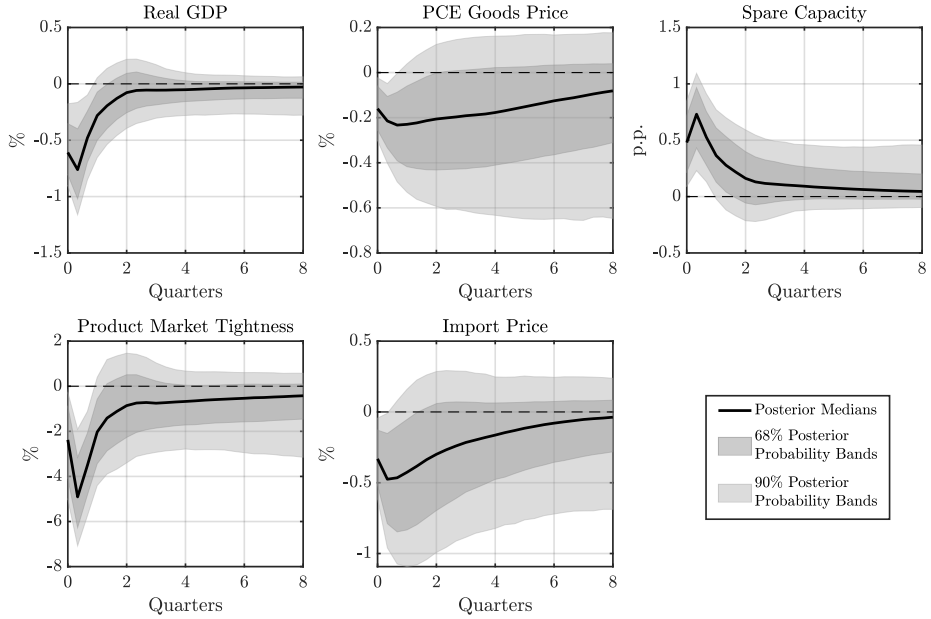

Figure E.6: IRFs to an Adverse Shock to Aggregate Demand: Dropping the ACR Index

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified with the ACR index omitted from the estimation, and with Restrictions 1″, 2″, and 3″ imposed on the contemporaneous IRFs. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

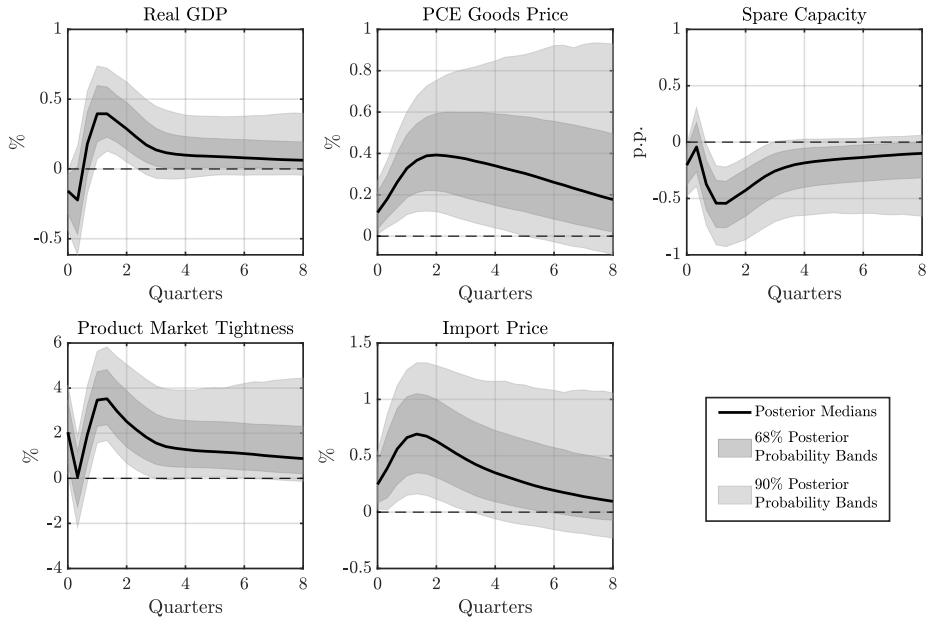

Figure E.7: IRFs to an Adverse Shock to Productive Capacity: Dropping the ACR Index

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified with the ACR index omitted from the estimation, and with Restrictions 1″, 2″, and 3″ imposed on the contemporaneous IRFs. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

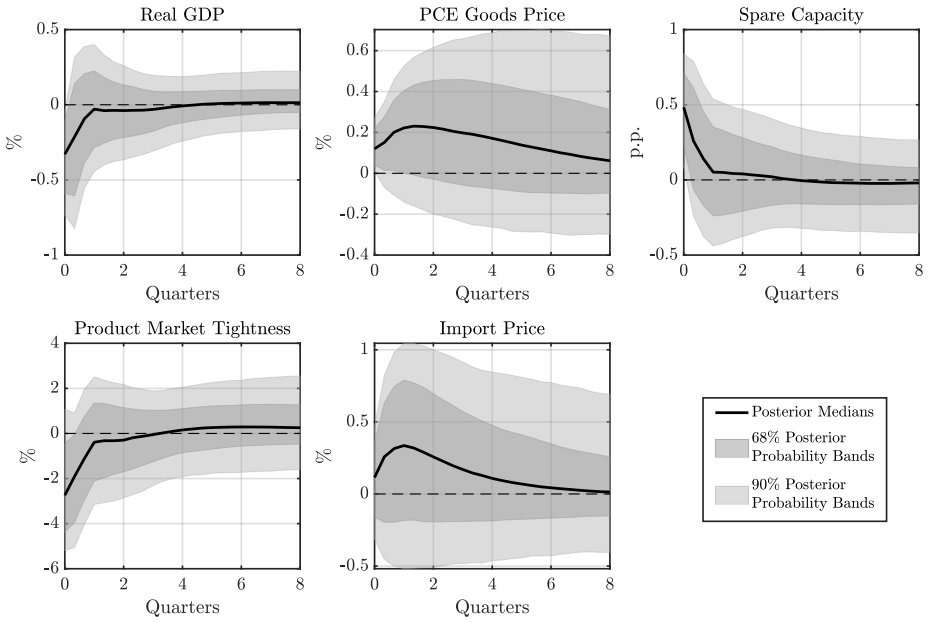

Figure E.8: IRFs to an Adverse Shock to the Supply Chain: Dropping the ACR Index

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified with the ACR index omitted from the estimation, and with Restrictions 1″, 2″, and 3″ imposed on the contemporaneous IRFs. The solid line shows the point-wise posterior medians, and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

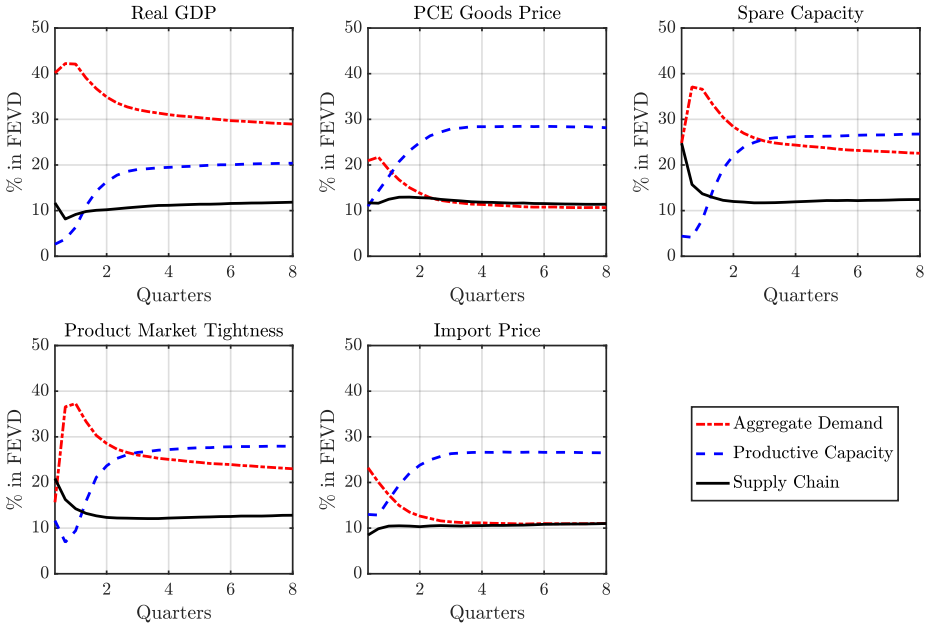

Figure E.9: FEVD from the SVAR: Dropping the ACR Index

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated with the ACR index omitted from the SVAR model, and with Restrictions 1″, 2″, and 3″ imposed on the contemporaneous IRFs. The figure is based on 100,000 independent importance sampling draws.

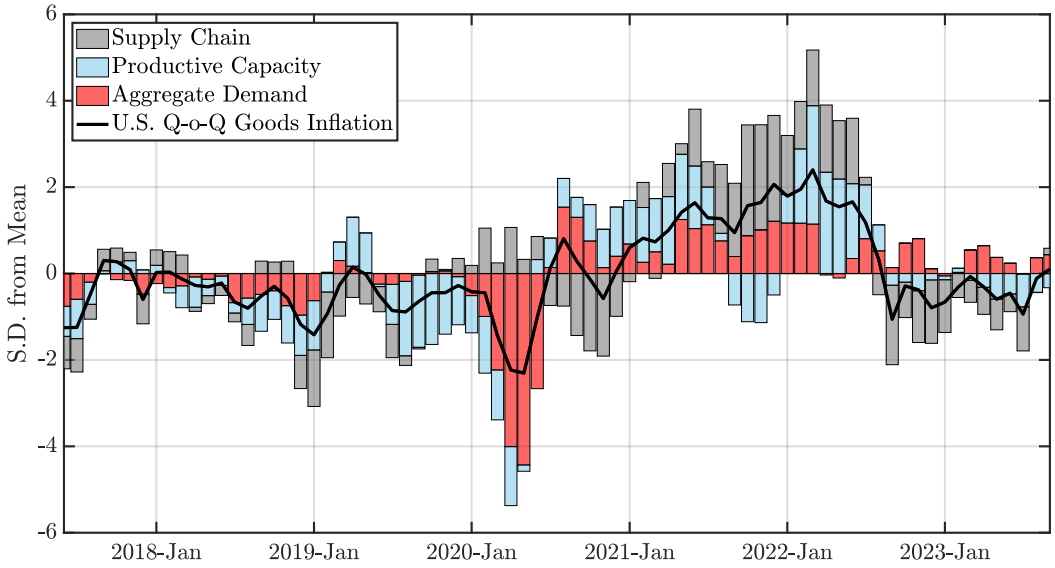

Figure E.10: HD of U.S. Quarter-on-Quarter Goods Inflation: Dropping the ACR Index

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with the ACR index omitted from the estimation, and with Restrictions $1''$, $2''$, and $3''$ imposed on the contemporaneous IRFs. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

E.3. Different Lag Structures

In the SVAR estimation, the choice of lag structure critically influences the IRFs since too few lags might miss important dynamics, leading to incomplete IRF patterns. At the same time, too many lags can cause over-fitting and spurious results. Therefore, conducting robustness checks by varying lag lengths and assessing the consistency of IRFs is key.

In Figures E.11 through E.19, we show that the IRFs of adverse shocks to aggregate demand, productive capacity, and the supply chain are robust to considering different lag structures, i.e., one, three, or four lags. We do not consider higher lags due to parameter uncertainty resulting from our limited sample length.

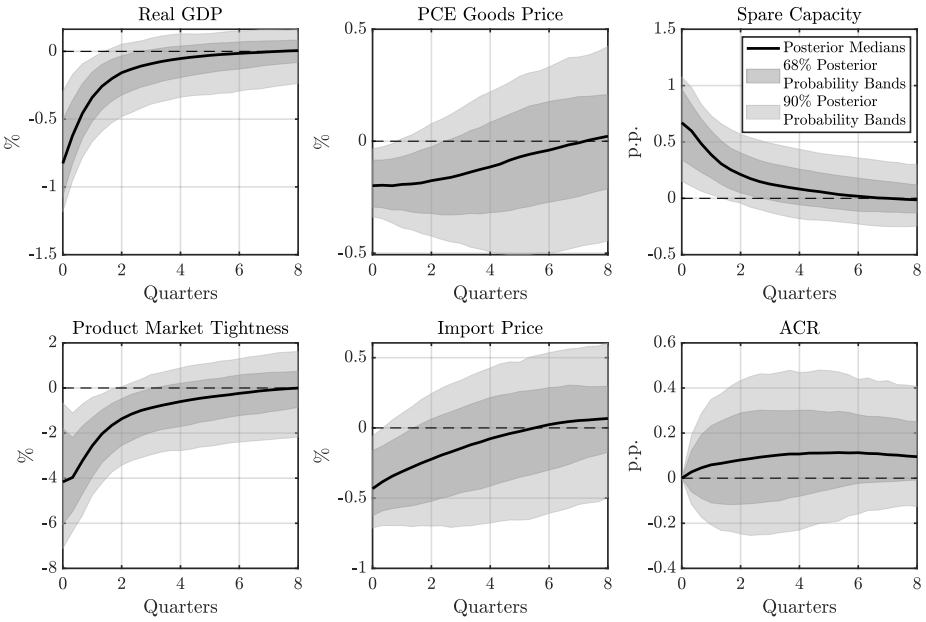

Figure E.11: IRFs to an Adverse Shock to Aggregate Demand: One Lag

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using an SVAR specification in Equation (28) with one lag. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

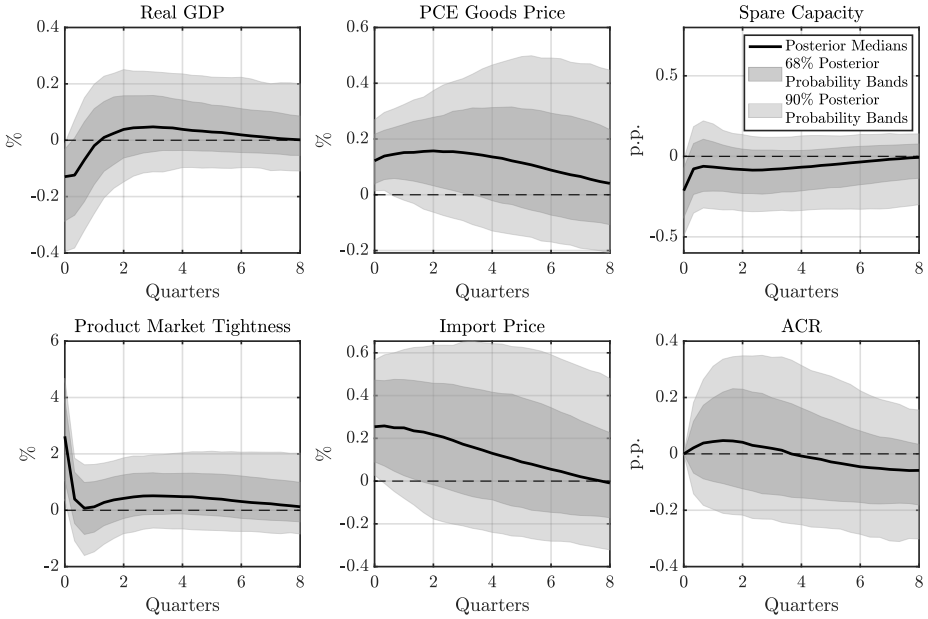

Figure E.12: IRFs to an Adverse Shock to Productive Capacity: One Lag

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using an SVAR specification in Equation (28) with one lag. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

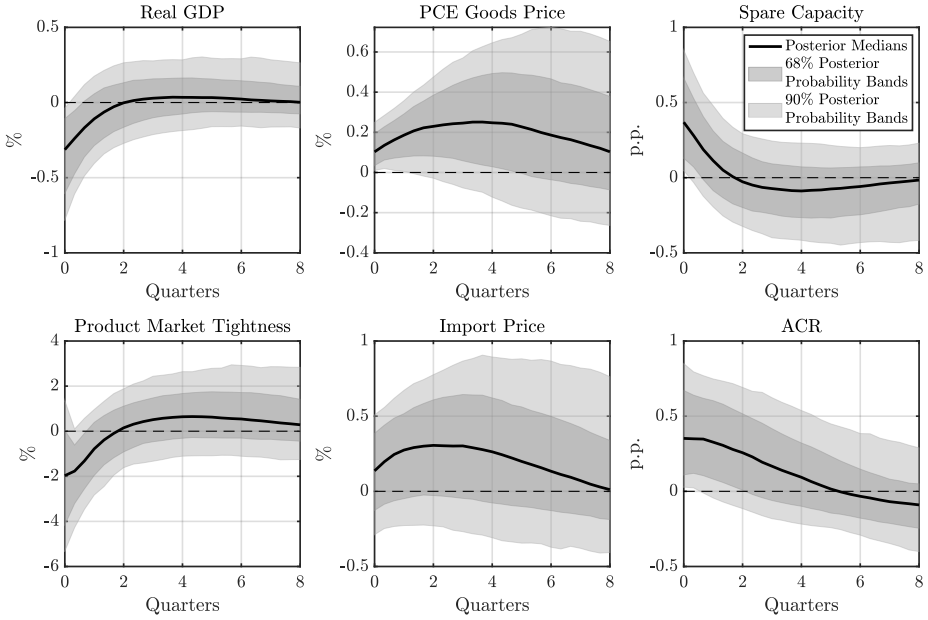

Figure E.13: IRFs to an Adverse Shock to the Supply Chain: One Lag

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using an SVAR specification in Equation (28) with one lag. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

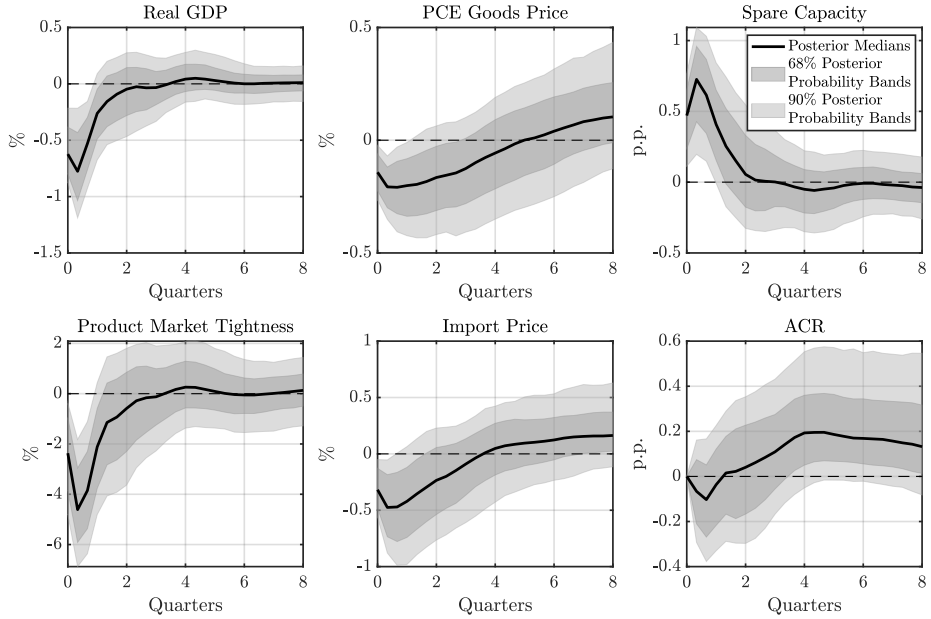

Figure E.14: IRFs to an Adverse Shock to Aggregate Demand: Three Lags

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using an SVAR specification in Equation (28) with three lags. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

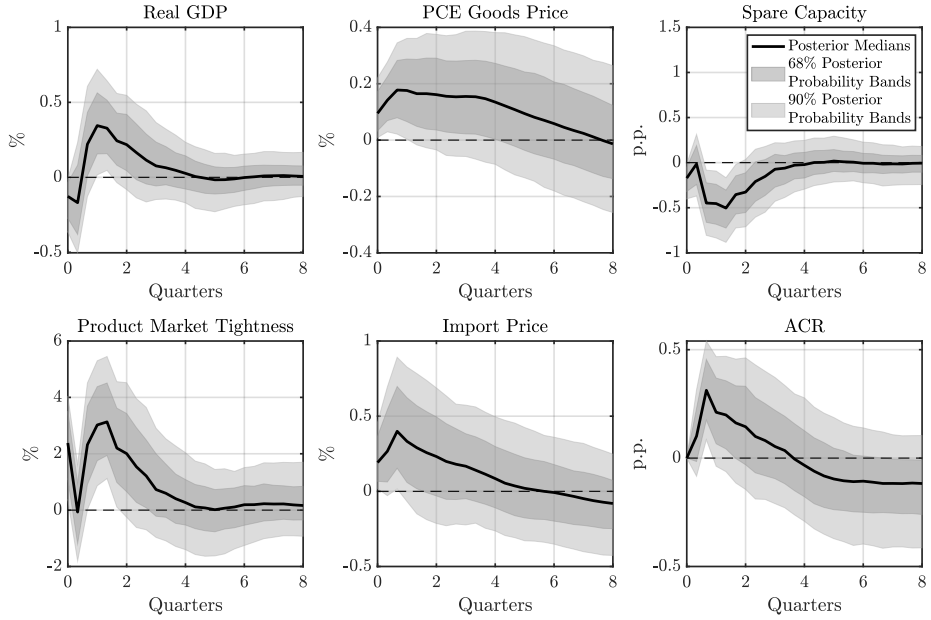

Figure E.15: IRFs to an Adverse Shock to Productive Capacity: Three Lags

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using an SVAR specification in Equation (28) with three lags. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

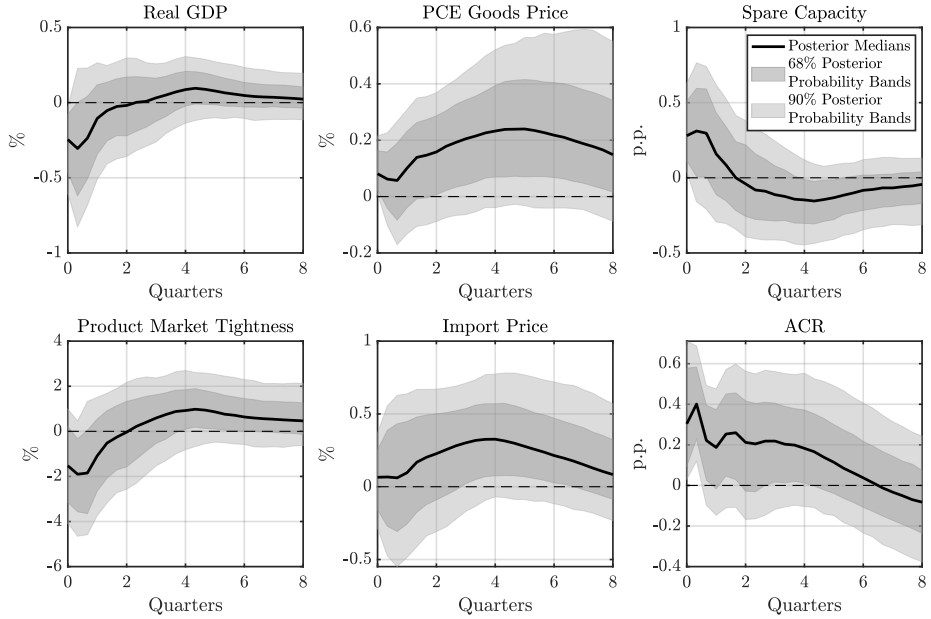

Figure E.16: IRFs to an Adverse Shock to the Supply Chain: Three Lags

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using an SVAR specification in Equation (28) with three lags. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

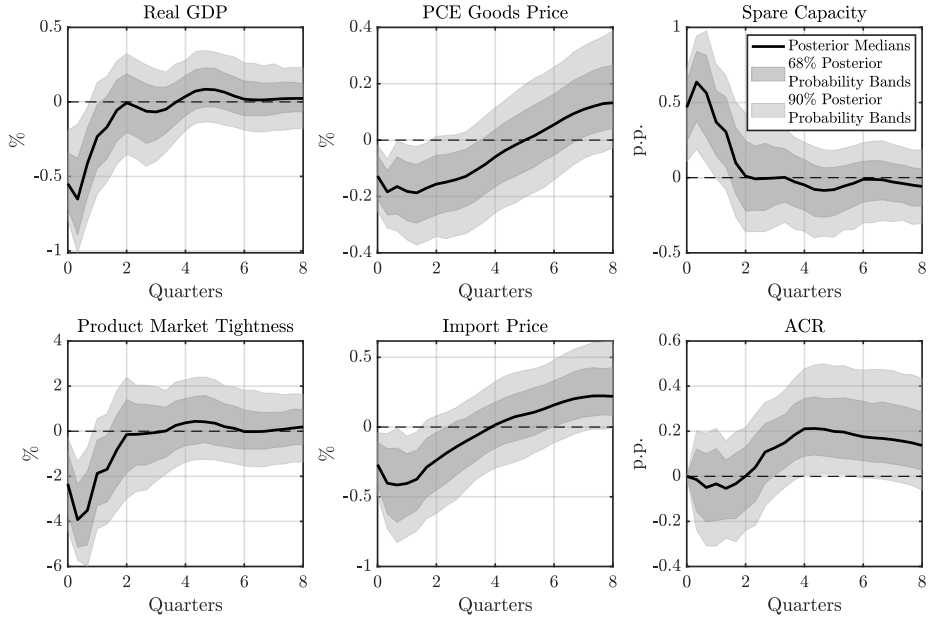

Figure E.17: IRFs to an Adverse Shock to Aggregate Demand: Four Lags

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using an SVAR specification in Equation (28) with four lags. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

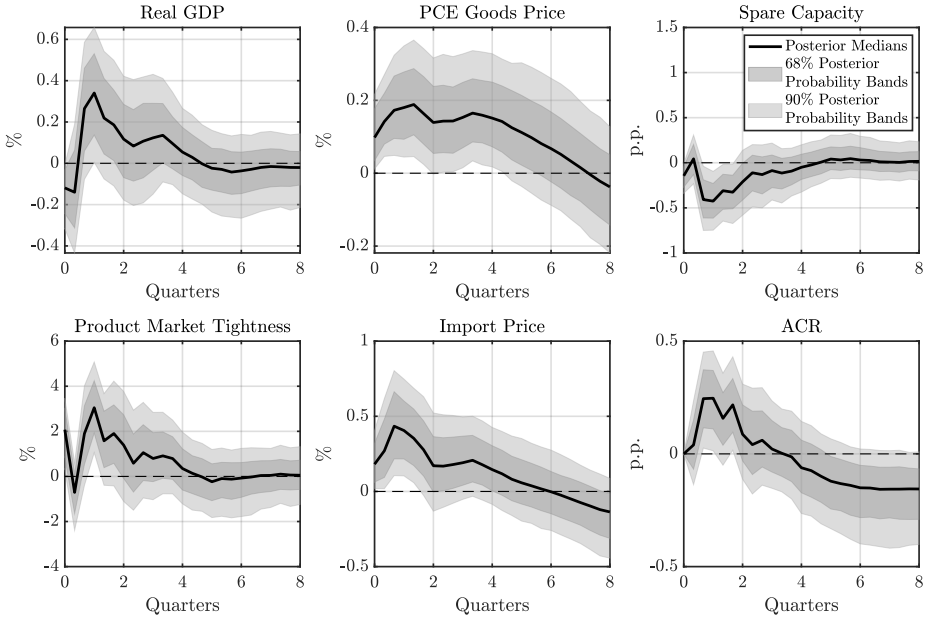

Figure E.18: IRFs to an Adverse Shock to Productive Capacity: Four Lags

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using an SVAR specification in Equation (28) with four lags. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

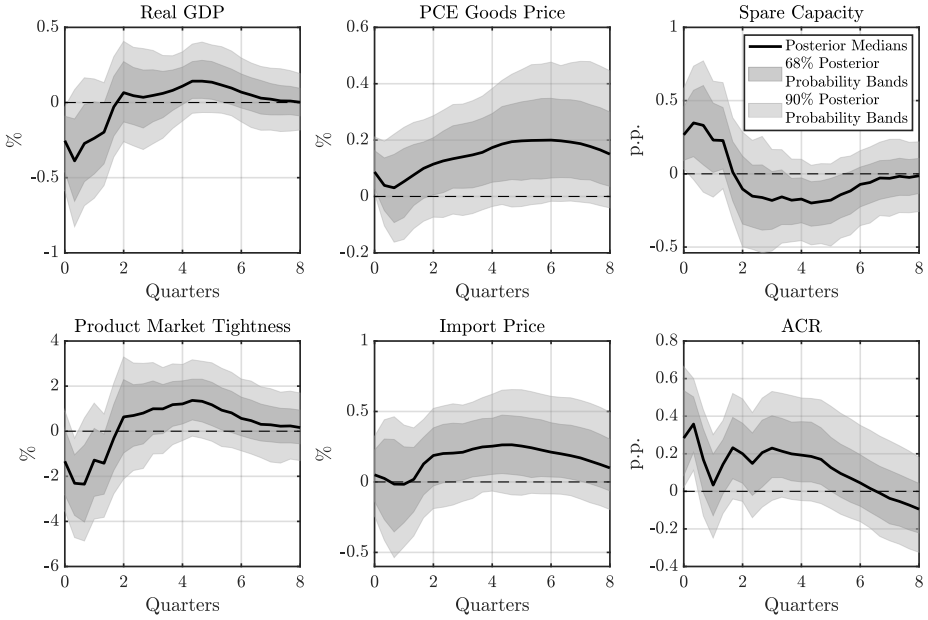

Figure E.19: IRFs to an Adverse Shock to the Supply Chain: Four Lags

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using an SVAR specification in Equation (28) with four lags. All other estimation specifications remain the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

E.4. Alternative Proxies for Consumption, Prices, and Product Market Tightness

In this appendix, we examine the robustness of our baseline results by replacing real GDP with the real PCE of goods and the industrial production index, the PCE goods price index with the GDP deflator, and the import price index with the producer price index in the SVAR estimation, respectively. The monthly time series for the real PCE of goods and the industrial production index are retrieved from the FRED database using the mnemonics `DGDSRX1` and `INDPRO`. The monthly time series for the GDP deflator is derived by interpolating the corresponding quarterly series using the Chow-Lin method (Chow and Lin, 1971), based on both the consumer price index and the producer price index, with the mnemonics `GDPDEF`, `CPIAUCSL`, and `WPSFD49207` in FRED. As shown in Figures E.20 through E.31, the IRFs in these alternative scenarios remain consistent with the baseline estimation.

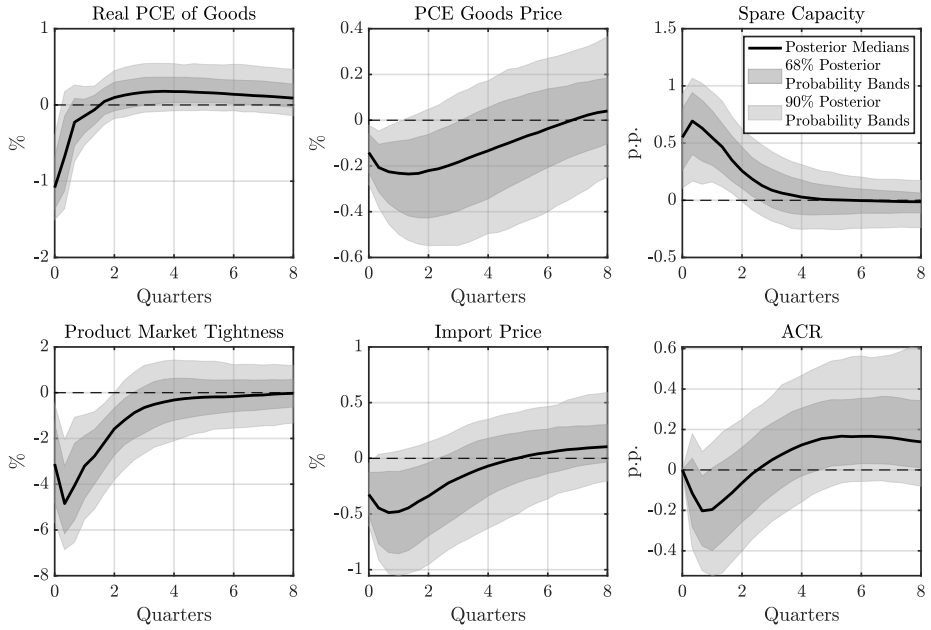

Figure E.20: IRFs to an Adverse Shock to Aggregate Demand: Real PCE of Goods

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using an SVAR specification in Equation (28), with the real PCE of goods serving as the proxy for consumption (output). All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

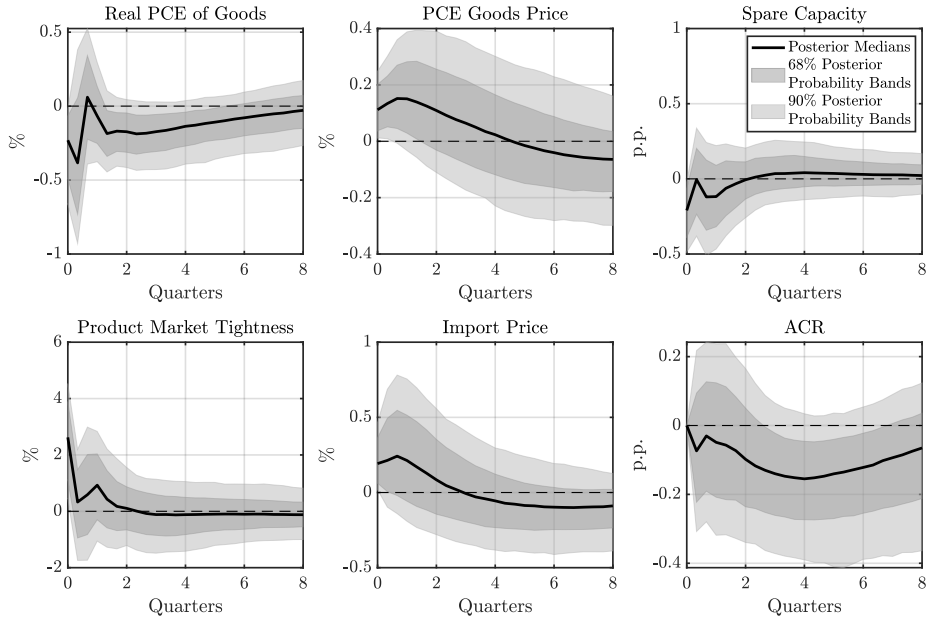

Figure E.21: IRFs to an Adverse Shock to Productive Capacity: Real PCE of Goods

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using an SVAR specification in Equation (28), with the real PCE of goods serving as the proxy for consumption (output). All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

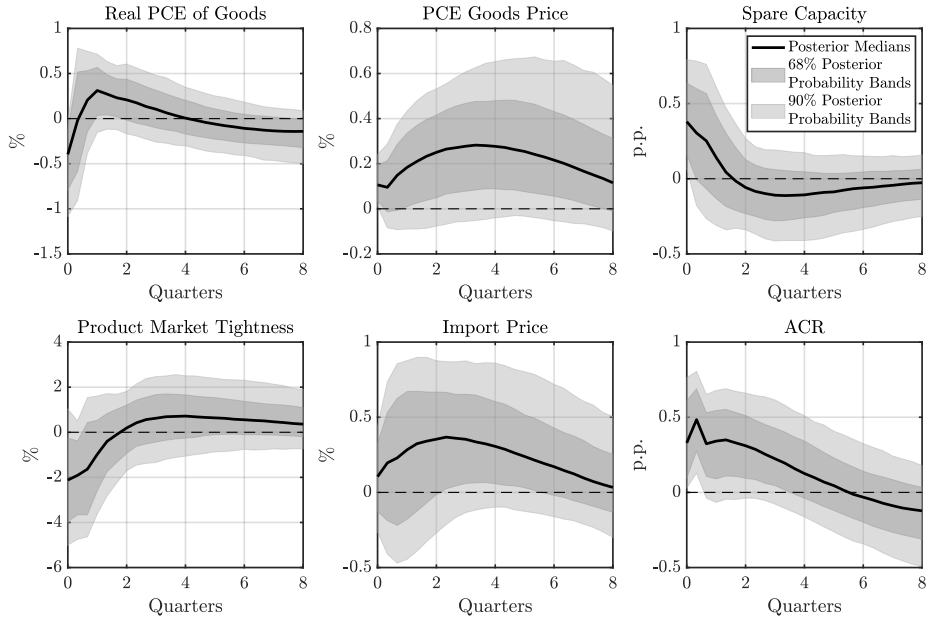

Figure E.22: IRFs to an Adverse Shock to the Supply Chain: Real PCE of Goods

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using an SVAR specification in Equation (28), with the real PCE of goods serving as the proxy for consumption (output). All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

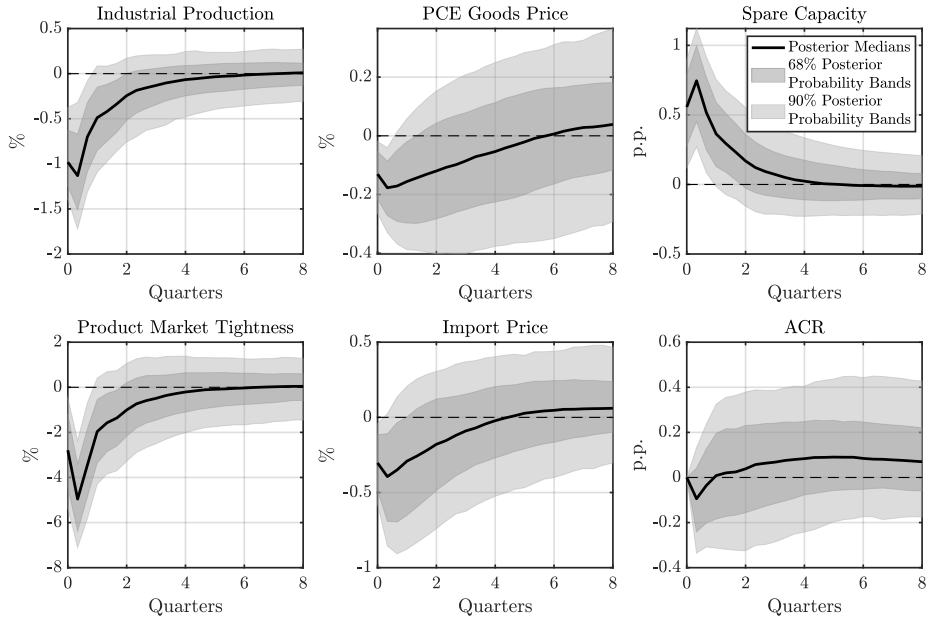

Figure E.23: IRFs to an Adverse Shock to Aggregate Demand: Industrial Production

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using an SVAR specification in Equation (28), with the industrial production index serving as the proxy for consumption (output). All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

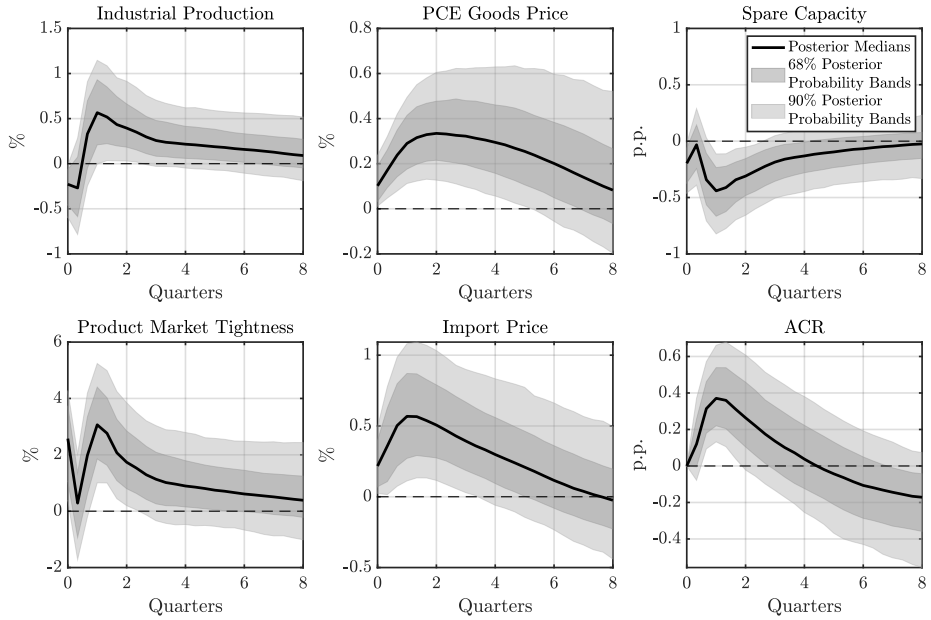

Figure E.24: IRFs to an Adverse Shock to Productive Capacity: Industrial Production

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using an SVAR specification in Equation (28), with the industrial production index serving as the proxy for consumption (output). All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

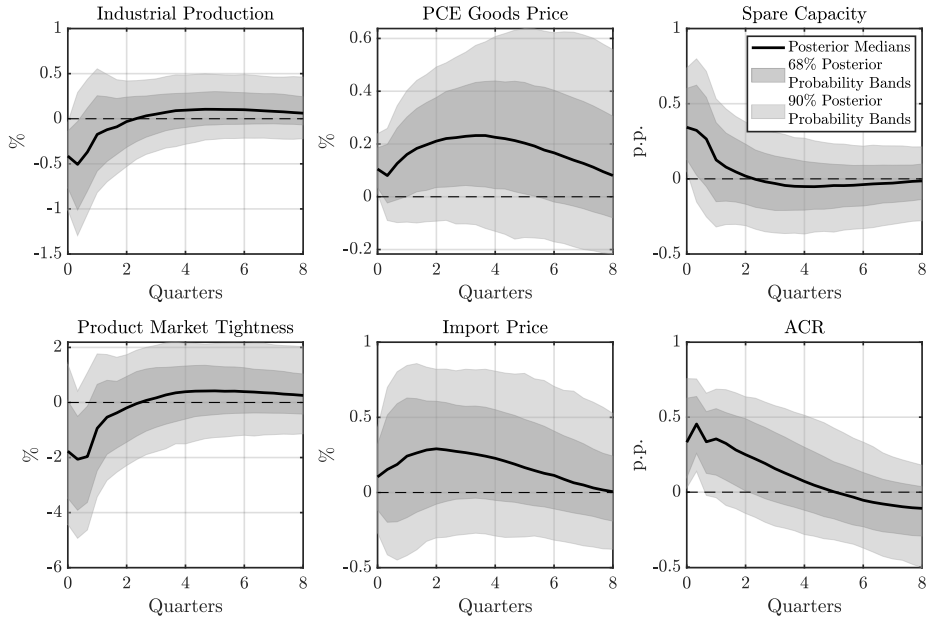

Figure E.25: IRFs to an Adverse Shock to the Supply Chain: Industrial Production

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using an SVAR specification in Equation (28), with the industrial production index serving as the proxy for consumption (output). All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

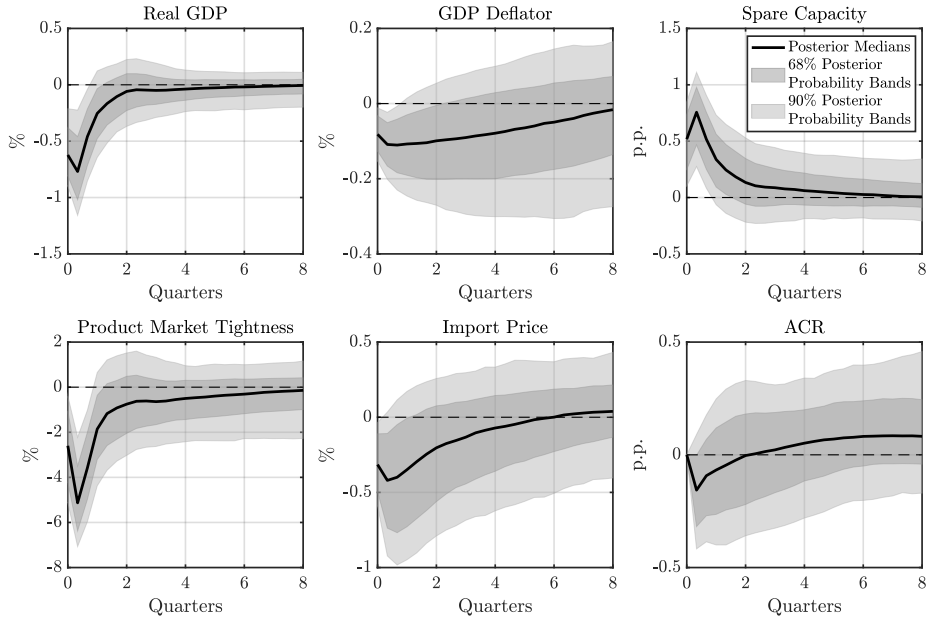

Figure E.26: IRFs to an Adverse Shock to Aggregate Demand: GDP Deflator

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using an SVAR specification in Equation (28), with the GDP deflator serving as the proxy for price. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

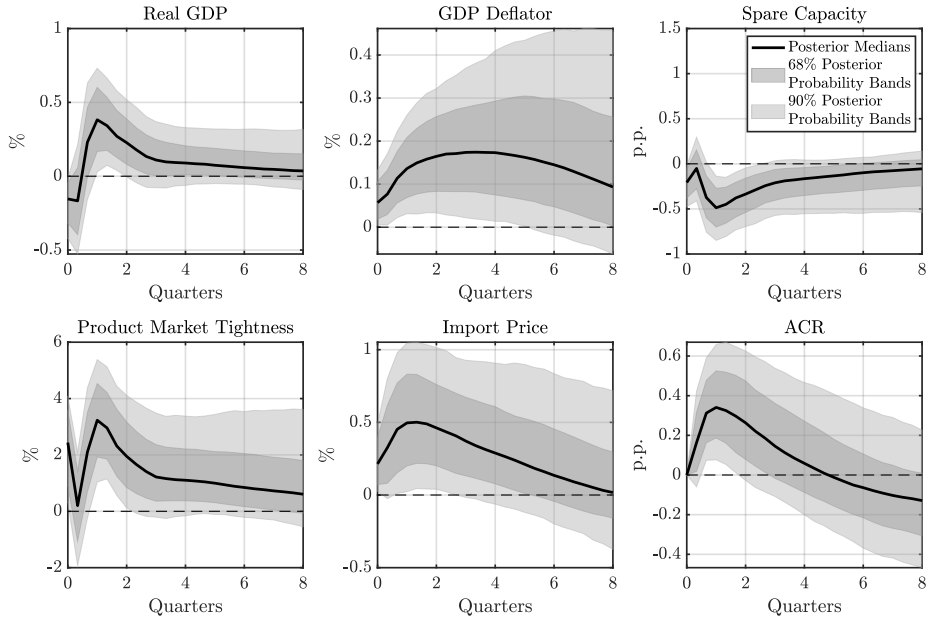

Figure E.27: IRFs to an Adverse Shock to Productive Capacity: GDP Deflator

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using an SVAR specification in Equation (28), with the GDP deflator serving as the proxy for price. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

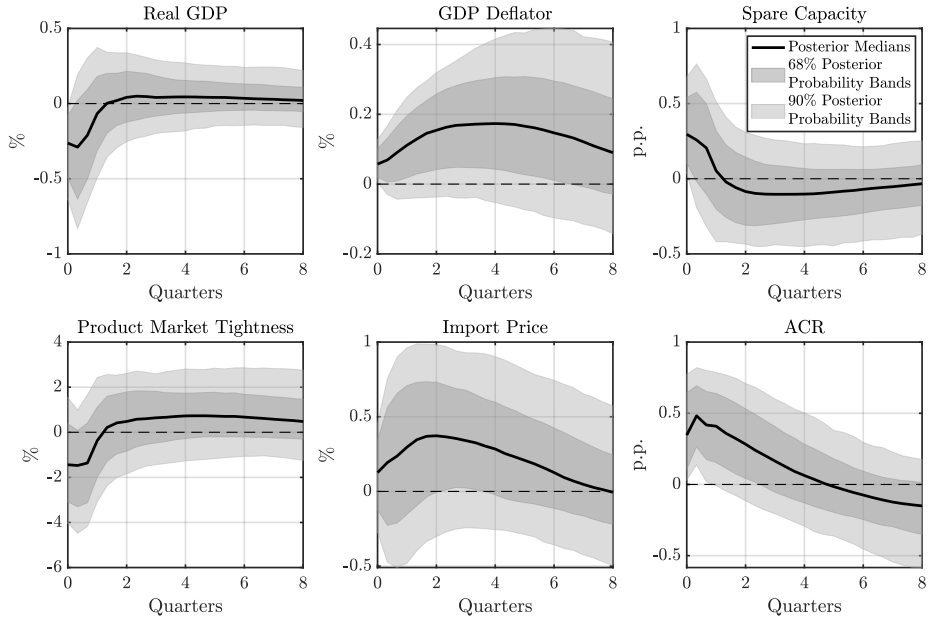

Figure E.28: IRFs to an Adverse Shock to the Supply Chain: GDP Deflator

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using an SVAR specification in Equation (28), with the GDP deflator serving as the proxy for price. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

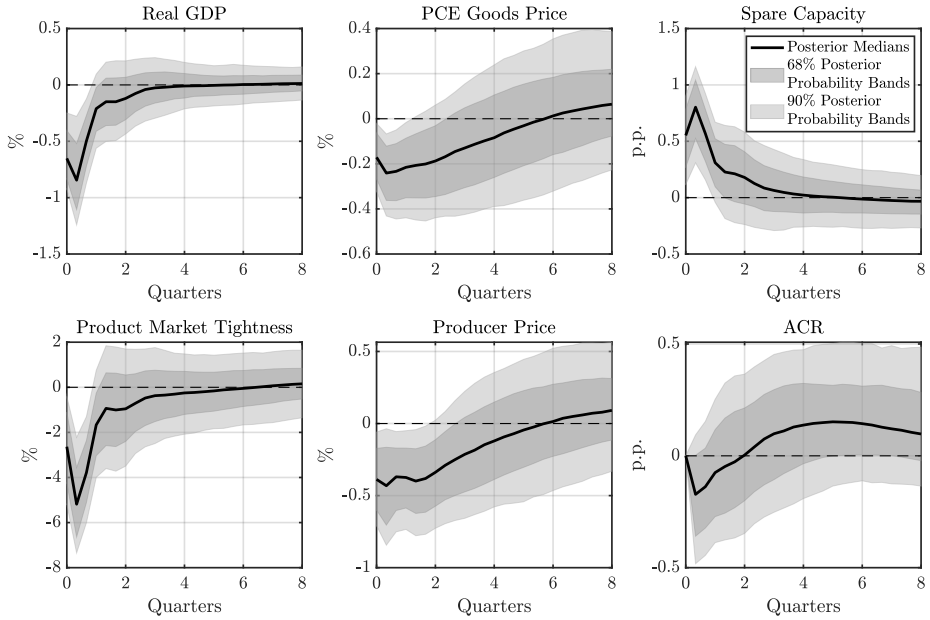

Figure E.29: IRFs to an Adverse Shock to Aggregate Demand: Producer Price

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using an SVAR specification in Equation (28), with the producer price index serving as the proxy for the wholesale price. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

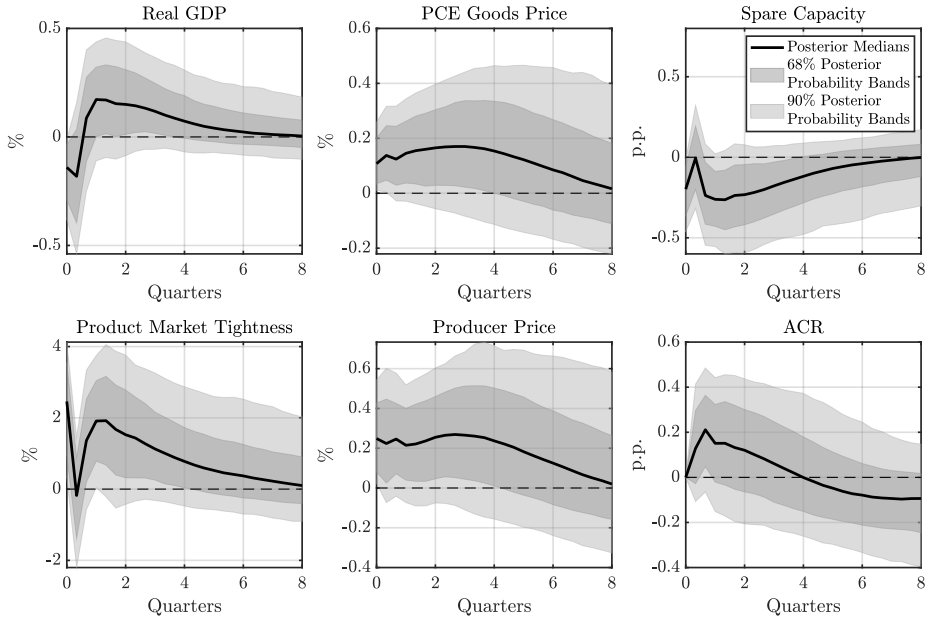

Figure E.30: IRFs to an Adverse Shock to Productive Capacity: Producer Price

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using an SVAR specification in Equation (28), with the producer price index serving as the proxy for the wholesale price. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

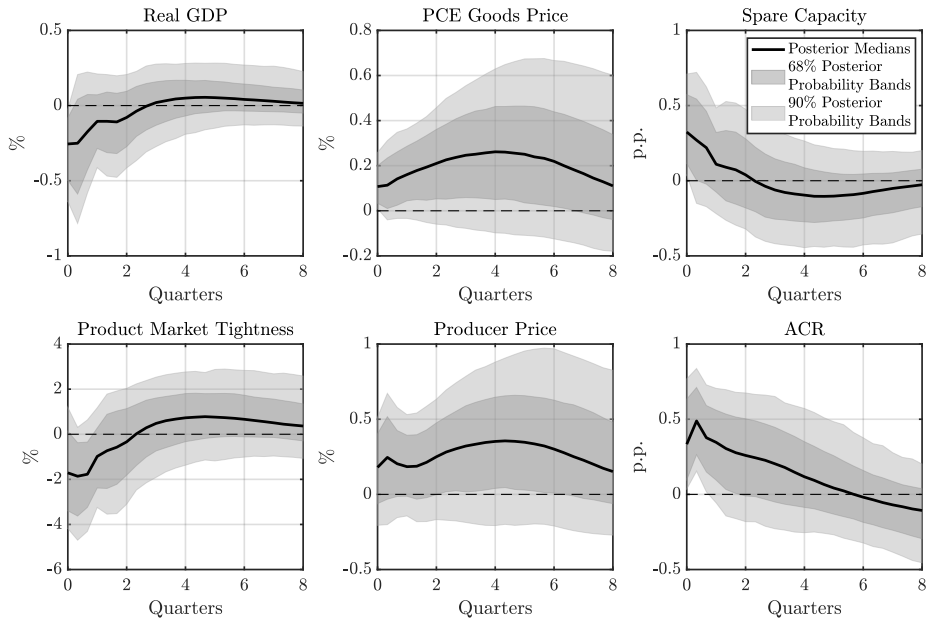

Figure E.31: IRFs to an Adverse Shock to the Supply Chain: Producer Price

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using an SVAR specification in Equation (28), with the producer price index serving as the proxy for the wholesale price. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

To further assess the robustness of our findings, we substitute the total U.S. manufacturers' new orders with the total U.S. retailers' new orders to construct an alternative measure of product market tightness. Specifically, we construct the monthly time series for this alternative measure by dividing the new orders placed by U.S. retailers by the import-weighted total spare capacity of the top five exporting countries to the U.S., namely, Mexico, Canada, China, Germany, and Japan, which collectively account for more than 53% of goods imported into the U.S.

For the retailers' new orders ($RetailNewOrder_t$), we estimate the corresponding U.S. dollar amount using the following formula:

$$\begin{aligned} RetailNewOrder_t &= \left(Inventory_t - Inventory_{t-1}\right) + Sale_t \\ &= \left(Inventory_t - Inventory_{t-1}\right) + \frac{Sales_t}{Inventory_t} \cdot Inventory_t, \end{aligned}$$

where $Inventory_t$ denotes the inventories held by U.S. retailers, and $Sale_t/Inventory_t$ represents the reciprocal of their inventories to sales ratio in month t. The monthly time series for U.S. retailers' inventories and their inventories to sales ratio can be retrieved directly from the FRED database using the mnemonics `RETAILIMSA` and `RETAILIRSA`.

For the import-weighted total spare capacity of the top five exporting countries to the U.S. ($SpareCapacityDollar_t$), we estimate its U.S. dollar amount using the following formula:

$$SpareCapacityDollar_t = \sum_{i \in \mathscr{C}} \left[\frac{Import_{i,t}}{\sum_{i \in \mathscr{C}} Import_{i,t}} \cdot \left(\frac{1}{CapacityUtilization_{i,t}} - 1 \right) \cdot IndustrialProduction_{i,t} \right],$$

where $\mathscr{C} \equiv \{Mexico, Canada, China, Germany, Japan\}$, $Import_{i,t}$ denotes the U.S. imports of goods by customs basis from country i in month t, $CapacityUtilization_{i,t}$ represents the capacity utilization rate for country i in month t, and $IndustrialProduction_{i,t}$ stands for the total industrial production of country i in month t, measured in U.S. dollars. The monthly time series for U.S. imports from each of the five exporting countries can be retrieved directly from FRED. The capacity utilization rates for Mexico and Japan are readily available from official sources. In contrast, for Canada, China, and Germany, these rates are derived by interpolating the corresponding quarterly series using the Chow-Lin method (Chow and Lin, 1971), based on each country's industrial production index. Finally, the monthly industrial production for each country, measured in U.S. dollars, is calculated using the corresponding monthly industrial production index and the total dollar value of goods produced in the manufacturing, mining, and utilities sectors in each country.

The behaviors of the IRFs, as shown in Figures E.32 through E.34, are quantitatively similar to those observed in Figures 10, 11, and 12.

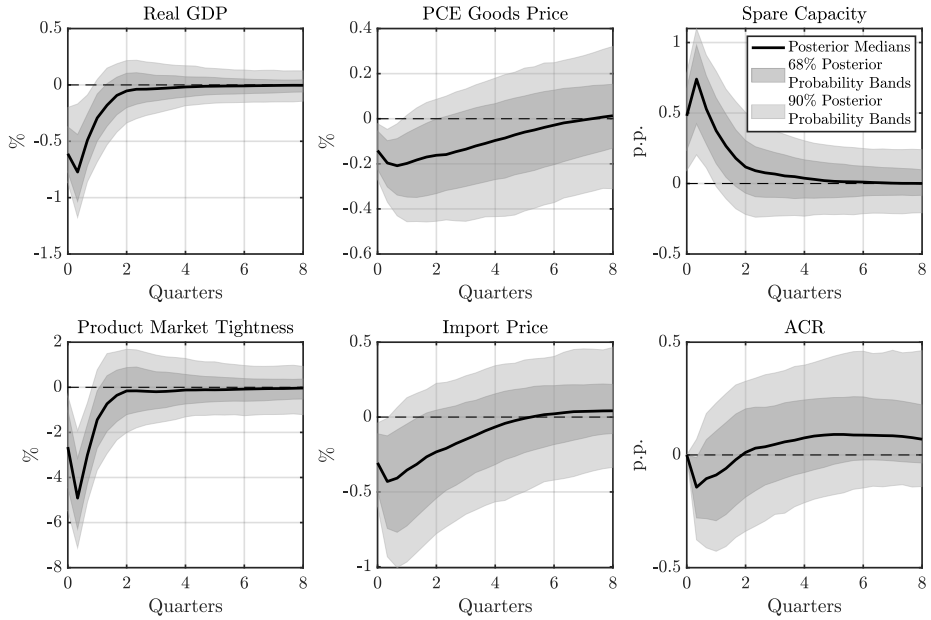

Figure E.32: IRFs to an Adverse Shock to Aggregate Demand: Retailers' New Orders

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified using an SVAR specification in Equation (28), with the total U.S. retailers' new orders serving as the proxy for the number of unmatched retailers in the construction of product market tightness. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

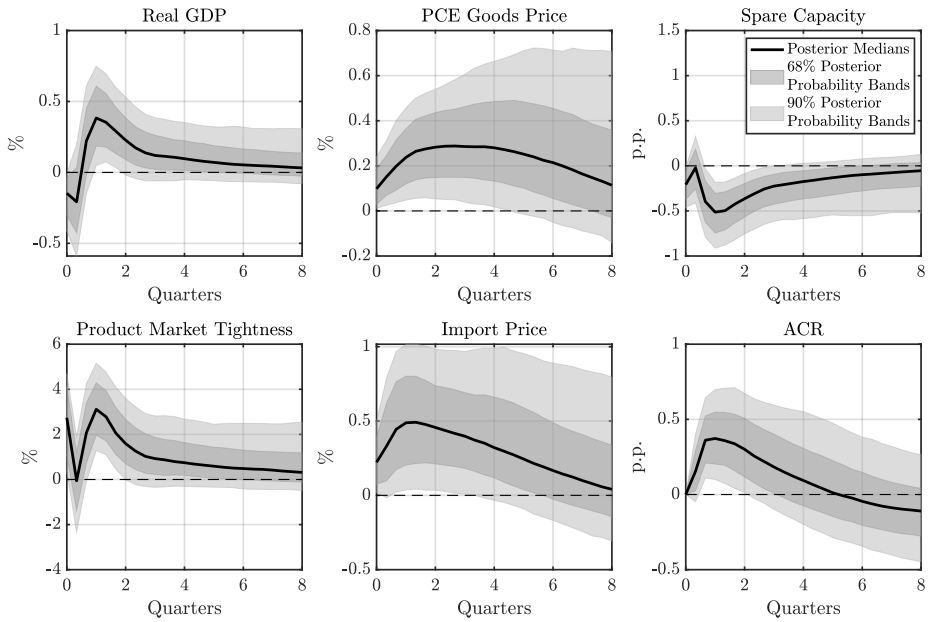

Figure E.33: IRFs to an Adverse Shock to Productive Capacity: Retailers' New Orders

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified using an SVAR specification in Equation (28), with the total U.S. retailers' new orders serving as the proxy for the number of unmatched retailers in the construction of product market tightness. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

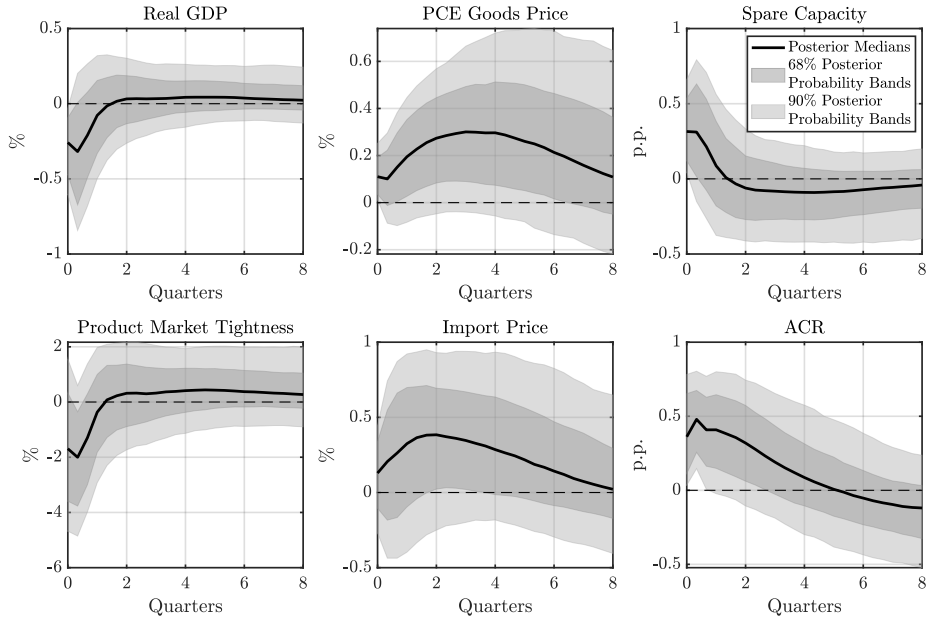

Figure E.34: IRFs to an Adverse Shock to the Supply Chain: Retailers' New Orders

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using an SVAR specification in Equation (28), with the total U.S. retailers' new orders serving as the proxy for the number of unmatched retailers in the construction of product market tightness. All other estimation specifications are kept the same as in the baseline. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

E.5. Prior Robustness

In this appendix, we demonstrate that the main conclusions from our baseline SVAR analysis remain robust when employing the prior robust approach for SVARs proposed by Giacomini and Kitagawa (2021). This approach eliminates the necessity of specifying a prior for the structural parameter based on the reduced-form parameter, which is primarily responsible for the asymptotic disagreement between Bayesian and frequentist inference. It is achieved by developing a class of priors that maintains a singular prior for the reduced-form parameter while permitting arbitrary conditional priors for the structural parameters, contingent upon the reduced-form parameter. This methodology enhances the robustness of our SVAR analysis, ensuring that our conclusions are not overly dependent on specific prior choices.

In practice, we apply their Algorithm 1 to numerically approximate the set of posterior means and the associated robust credible regions for the IRFs of the selected endogenous variables in response to each structural shock. We make two modifications in the implementation of Algorithm 1. First, in Step 2 of Algorithm 1, to draw the orthonormal Q's subject to Restrictions 1, 2, and 3, we apply the QR decomposition method as in Arias et al. (2018) instead of the original linear projection approach. These two ways of drawing Q's are comparable in terms of both the resulting distribution of Q and the computational cost. Second, we replace Step 3 of Algorithm 1 with Step 3′ of Algorithm 2 to approximate the lower and upper bounds of the prior robust posterior means, as well as those associated with the robust credible regions. These modifications enhance the precision and applicability of the algorithm in our context, ensuring a more robust and accurate approximation of the posterior means and credible regions for the IRFs.

In Figures E.35, E.36, and E.37, the solid lines represent the point-wise posterior medians. At the same time, the shaded areas depict the 68% equal-tailed point-wise posterior probability bands. These are based on the baseline estimation data from Section 4 in the main text. In addition, dotted curves are used to plot the set of prior robust posterior means, and the corresponding 68% robust credible regions are indicated with dashed-dotted curves. The underlying data for these plots are derived from 1,000 independent draws of the reduced-form parameters and 100,000 orthogonal matrix draws for each reduced-form parameter.

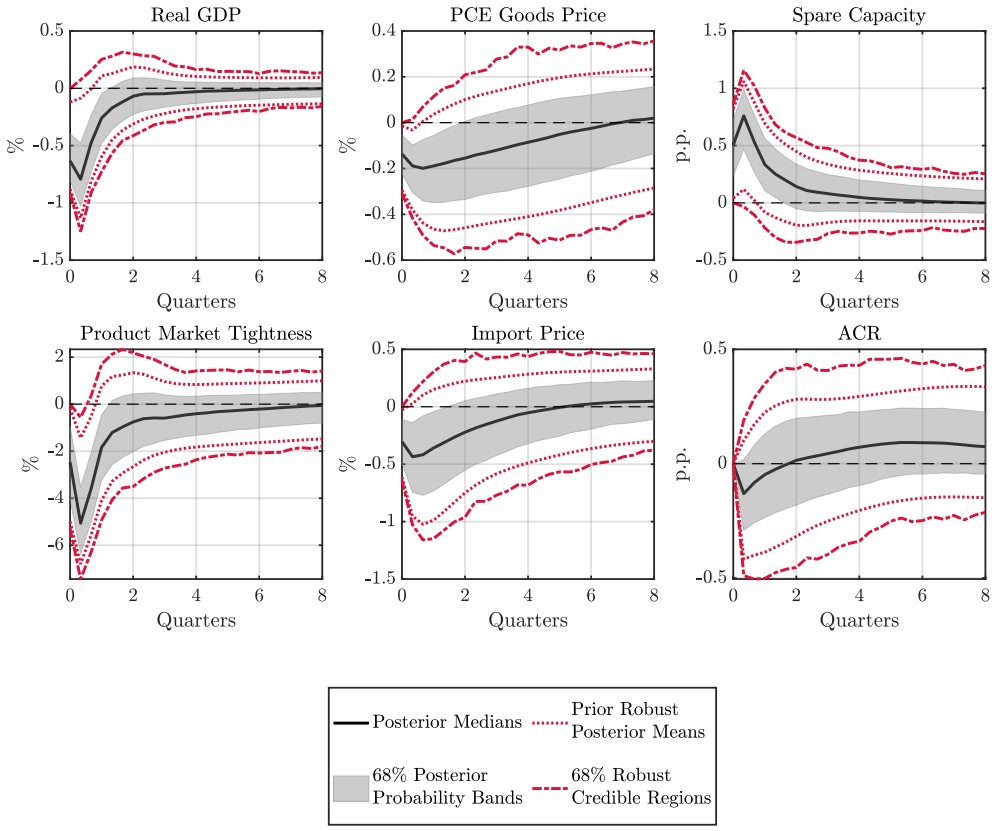

Figure E.35: IRFs to an Adverse Shock to Aggregate Demand: Prior Robustness

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are estimated using the prior robust approach for the SVARs proposed by Giacomini and Kitagawa (2021). The solid lines show the point-wise posterior medians and the shaded areas represent the 68% equal-tailed point-wise posterior probability bands, which are based on the data from our baseline estimation outlined in Section 4. The dotted curves illustrate the set of prior robust posterior means, and the dashed-dotted curves depict the 68% robust credible regions. These curves are obtained from 1,000 independent draws of the reduced-form parameters, and 100,000 orthogonal matrix draws for each reduced-form parameter.

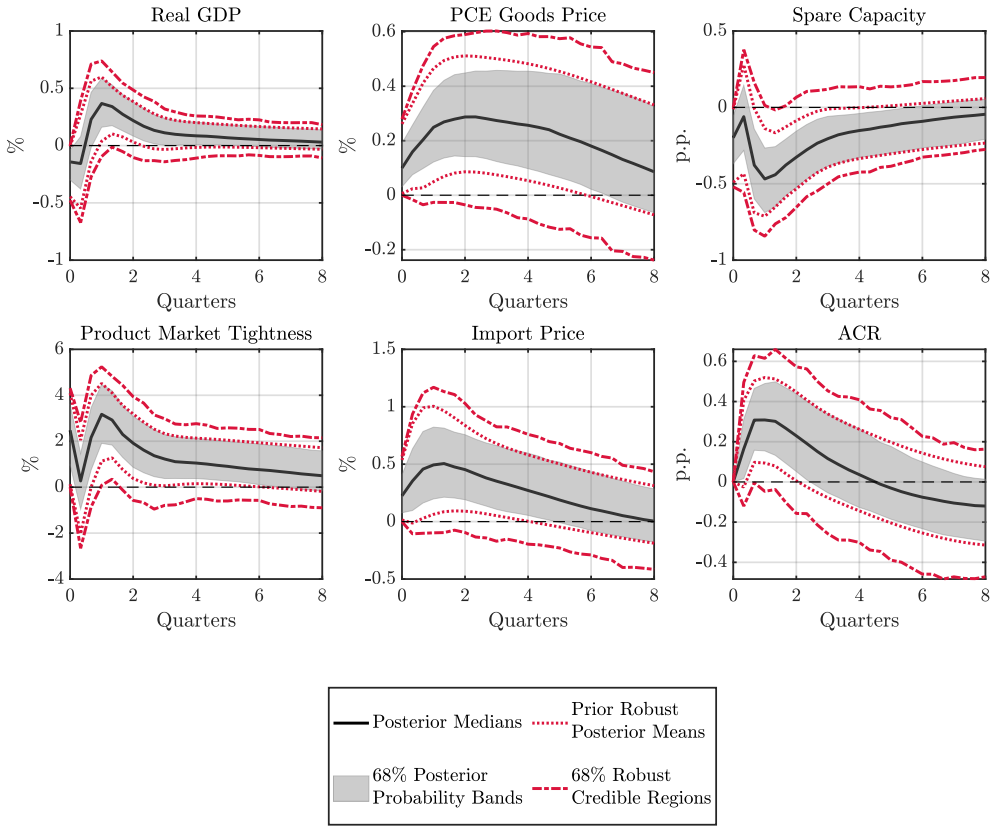

Figure E.36: IRFs to an Adverse Shock to Productive Capacity: Prior Robustness

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are estimated using the prior robust approach for the SVARs proposed by Giacomini and Kitagawa (2021). The solid lines show the point-wise posterior medians and the shaded areas represent the 68% equal-tailed point-wise posterior probability bands, which are based on the data from our baseline estimation outlined in Section 4. The dotted curves illustrate the set of prior robust posterior means, and the dashed-dotted curves depict the 68% robust credible regions. These curves are obtained from 1,000 independent draws of the reduced-form parameters, and 100,000 orthogonal matrix draws for each reduced-form parameter.

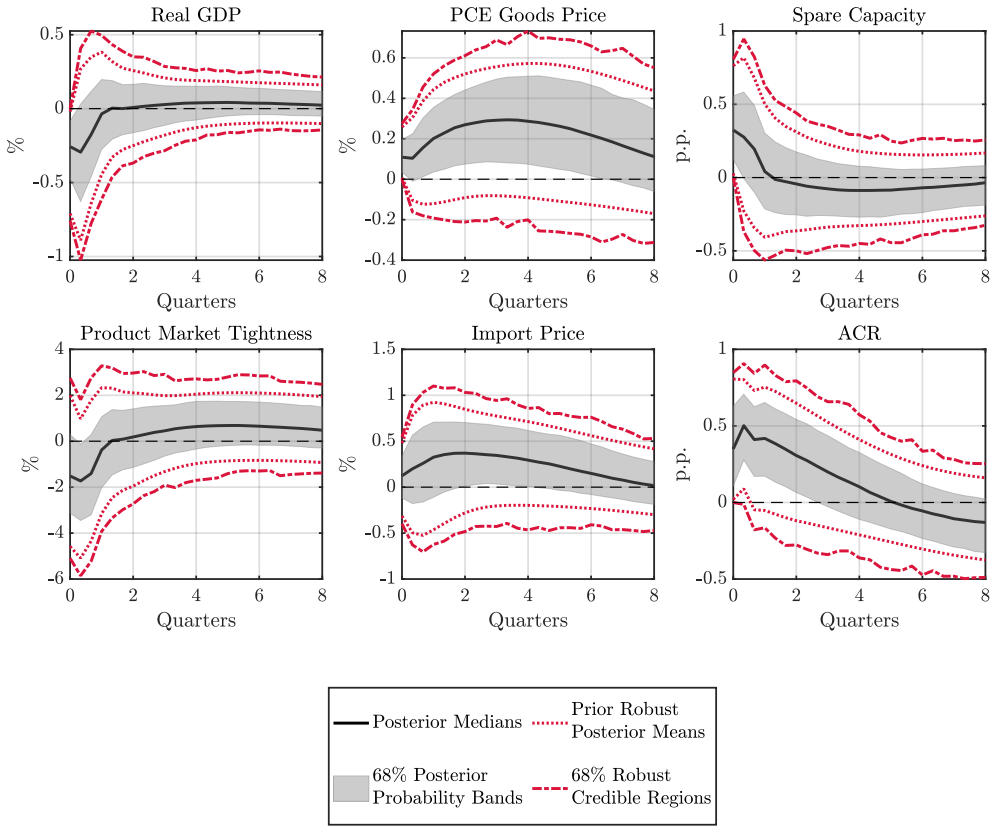

Figure E.37: IRFs to an Adverse Shock to the Supply Chain: Prior Robustness

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are estimated using the prior robust approach for the SVARs proposed by Giacomini and Kitagawa (2021). The solid lines show the point-wise posterior medians and the shaded areas represent the 68% equal-tailed point-wise posterior probability bands, which are based on the data from our baseline estimation outlined in Section 4. The dotted curves illustrate the set of prior robust posterior means, and the dashed-dotted curves depict the 68% robust credible regions. These curves are obtained from 1,000 independent draws of the reduced-form parameters, and 100,000 orthogonal matrix draws for each reduced-form parameter.

E.6. Additional Sign Restrictions on Product Market Tightness and Import Price

In an alternative SVAR estimation, together with Restrictions 1, 2, and 3, we apply additional positive sign restrictions on the contemporaneous responses of product market tightness and import price in response to an adverse supply chain shock. This robustness check aligns with our theoretical prediction in Proposition 5 for a supply chain disturbance represented by a reduction in matching efficiency. As illustrated in Figures E.38 through E.42, the results are quantitatively similar to those obtained without these additional sign restrictions.

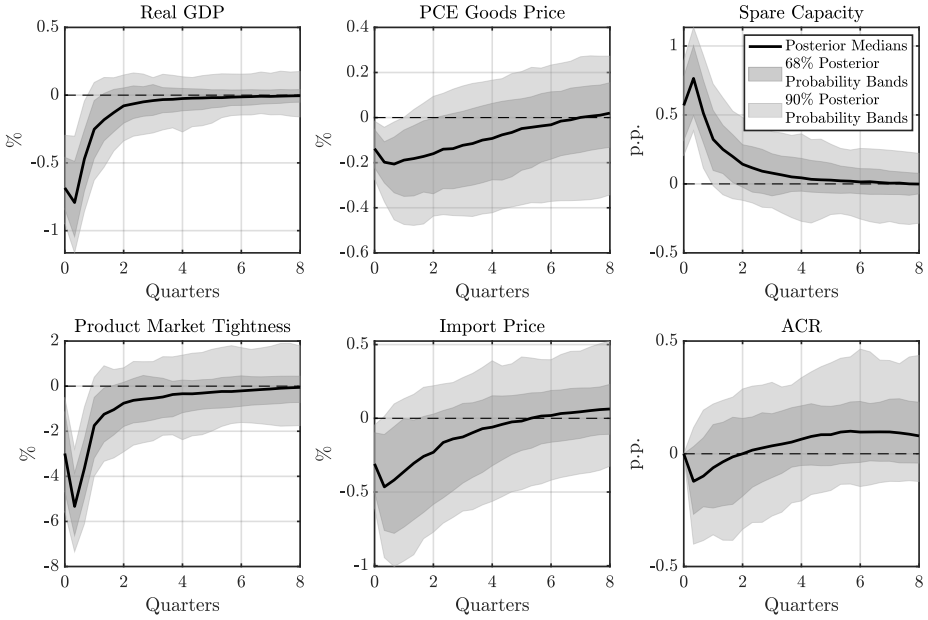

Figure E.38: IRFs to an Adverse Shock to Aggregate Demand: Additional Sign Restrictions

Notes. The IRFs to a one standard deviation adverse shock to aggregate demand are identified by imposing positive sign restrictions on the contemporaneous responses of product market tightness and import prices in the case of a supply chain disturbance, in conjunction with Restrictions 1, 2, and 3. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

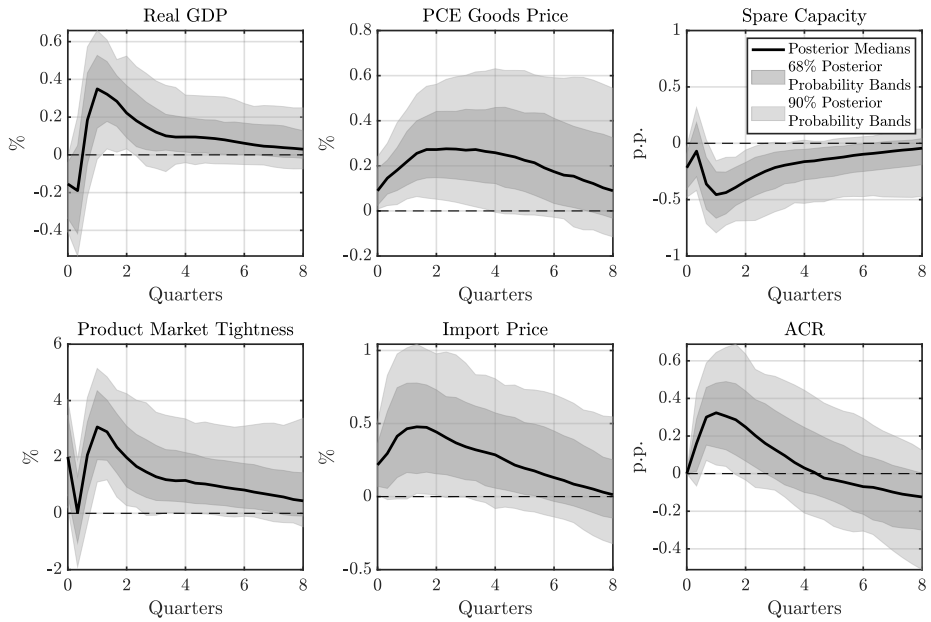

Figure E.39: IRFs to an Adverse Shock to Productive Capacity: Additional Sign Restrictions

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified by imposing positive sign restrictions on the contemporaneous responses of product market tightness and import prices in the case of a supply chain disturbance, in conjunction with Restrictions 1, 2, and 3. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

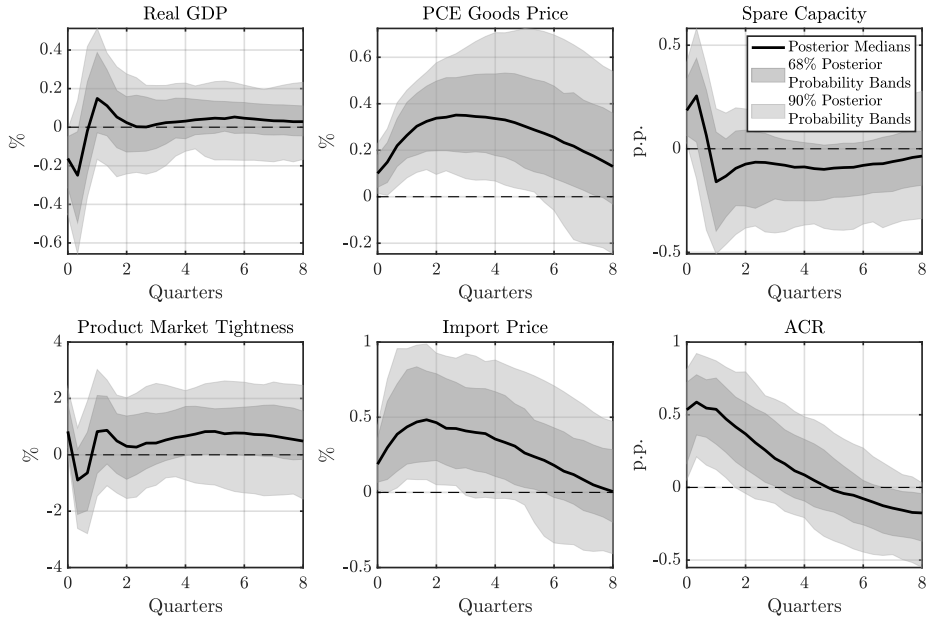

Figure E.40: IRFs to an Adverse Shock to the Supply Chain: Additional Sign Restrictions

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified by imposing positive sign restrictions on the contemporaneous responses of product market tightness and import prices in the case of a supply chain disturbance, in conjunction with Restrictions 1, 2, and 3. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

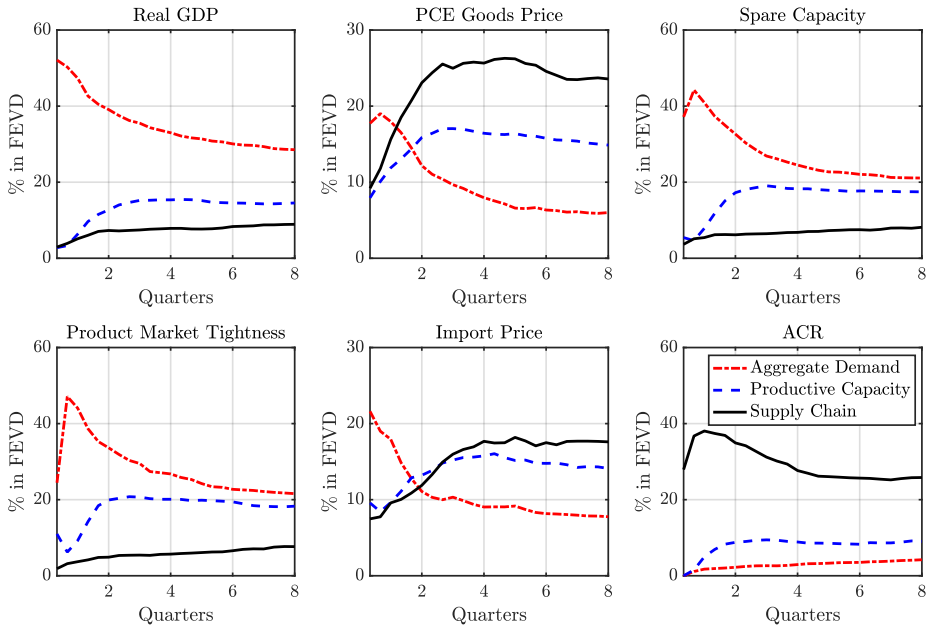

Figure E.41: FEVD from the SVAR: Additional Sign Restrictions

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated by applying positive sign restrictions on the contemporaneous responses of product market tightness and import price to a supply chain disturbance, in conjunction with Restrictions 1, 2, and 3. The figure is based on 100,000 independent importance sampling draws.

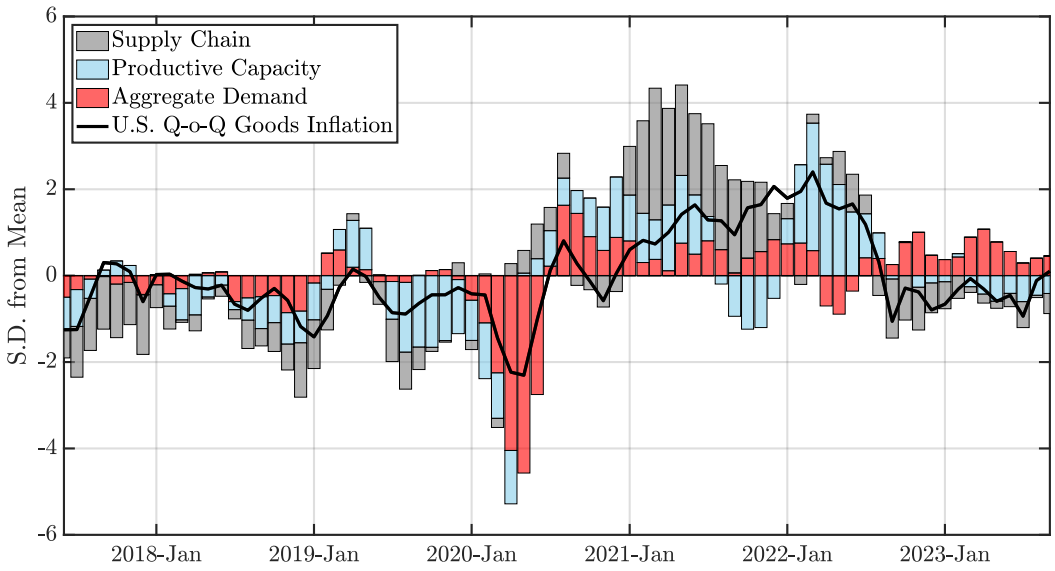

Figure E.42: HD of U.S. Quarter-on-Quarter Goods Inflation: Additional Sign Restrictions

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with the ACR index included as the measure indicating the state of the global supply chain. Positive sign restrictions are imposed on the contemporaneous responses of product market tightness and import prices to a supply chain disturbance, in conjunction with Restrictions 1, 2, and 3. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

E.7. Adding the Federal Funds Rate

In this appendix, we check whether the identification results for the supply chain disturbances are preserved by adding the federal funds rate to the SVAR estimation and by considering a contractionary monetary policy shock specifically. This involves imposing Restriction 4 alongside the following identification restrictions on the contemporaneous responses of the endogenous variables:

Restriction 2‴. ***An adverse shock to productive capacity*** *leads to a negative response of real GDP and spare capacity, as well as to a positive response of PCE goods price, product market tightness, and import price at* $k = 1$*. Both the ACR and the federal funds rate do not respond at* $k = 1$*.*

Restriction 3‴. ***An adverse shock to the supply chain*** *leads to a negative response of real GDP, as well as to a positive response of PCE goods price, spare capacity, and the ACR at* $k = 1$*. The federal funds rate does not respond at* $k = 1$*.*

To reduce parameter uncertainty, we include only one lag when estimating the SVAR model. The following figures show the estimated IRFs. Notably, as shown in Figure E.45, the stagflationary effect of a supply chain disruption shock is still observable, along with an accompanying jump in spare capacity on impact. However, the confidence bands are significantly wider than those shown in Figure 12.

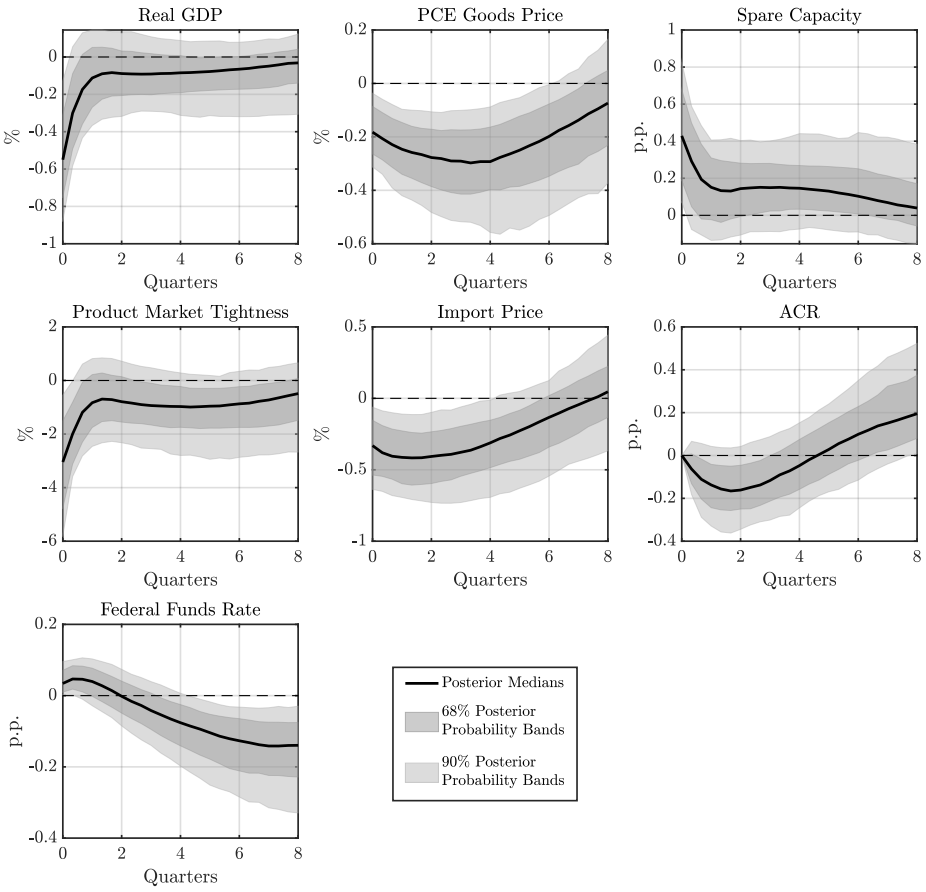

Figure E.43: IRFs to a Contractionary Monetary Policy Shock: Adding the Federal Funds Rate

Notes. The IRFs to a one standard deviation contractionary monetary policy shock are identified with the federal funds rate included in the SVAR model as an additional endogenous variable, and with Restrictions 2‴, 3‴, and 4 imposed on the contemporaneous IRFs. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

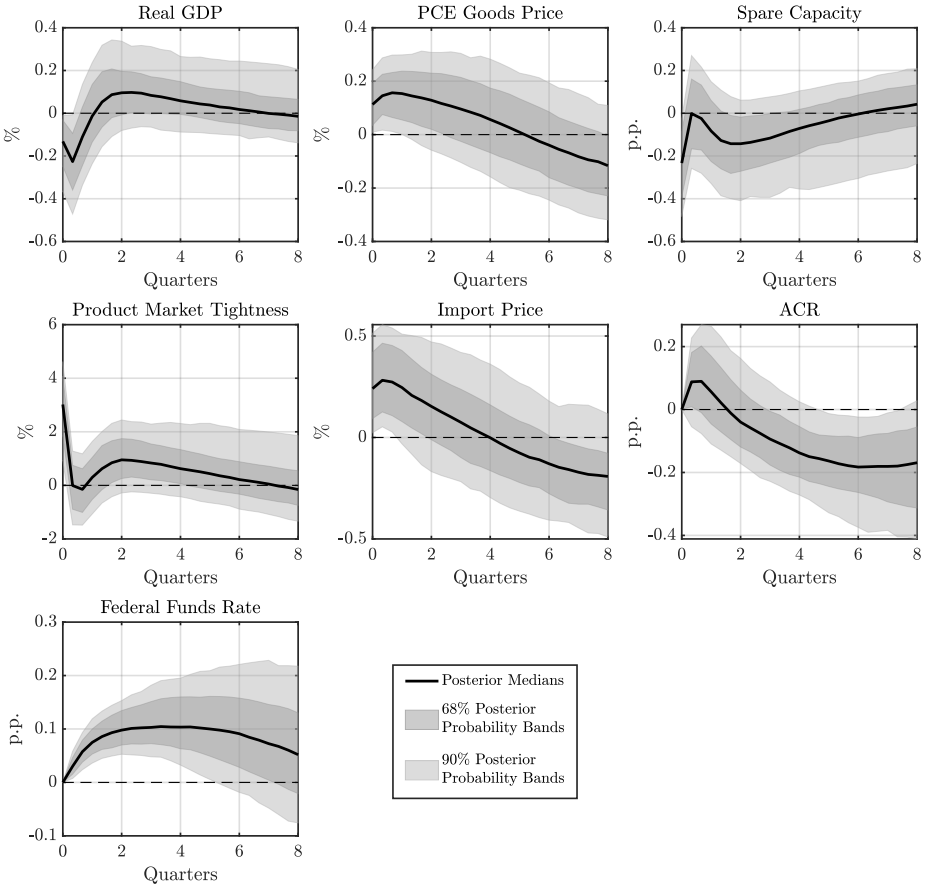

Figure E.44: IRFs to an Adverse Shock to Productive Capacity: Adding the Federal Funds Rate

Notes. The IRFs to a one standard deviation adverse shock to productive capacity are identified with the federal funds rate included in the SVAR model as an additional endogenous variable, and with Restrictions 2‴, 3‴, and 4 imposed on the contemporaneous IRFs. The solid line shows the point-wise posterior medians and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

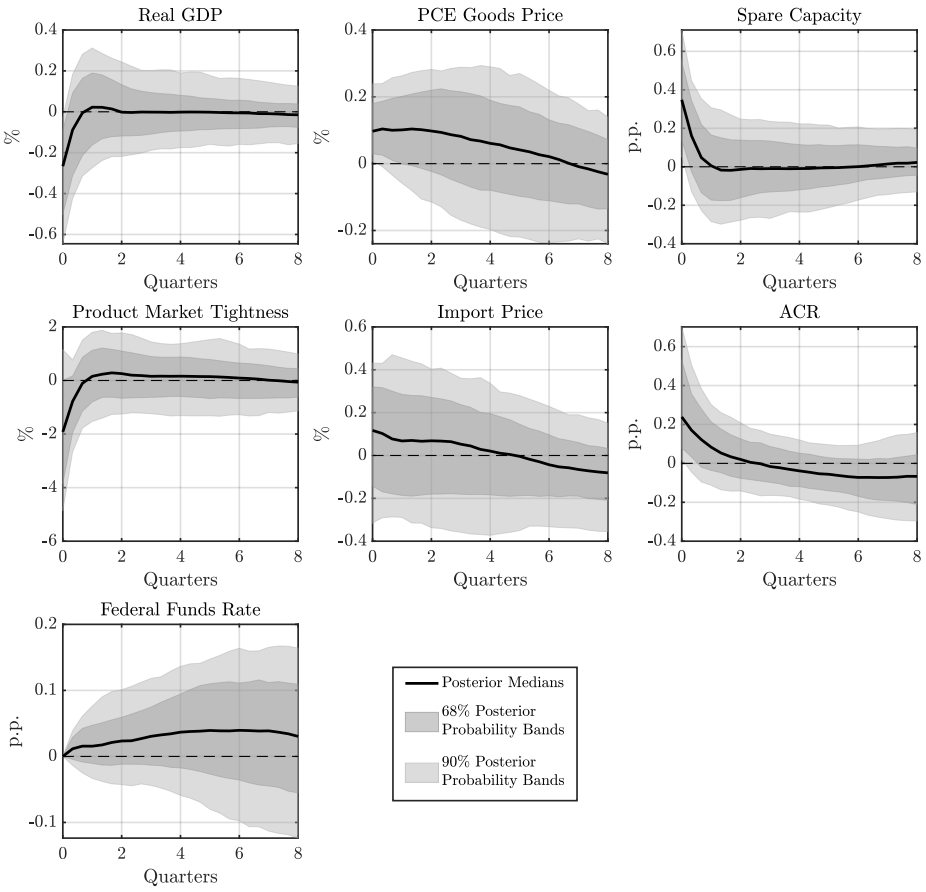

Figure E.45: IRFs to an Adverse Shock to the Supply Chain: Adding the Federal Funds Rate

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified with the federal funds rate included in the SVAR model as an additional endogenous variable, and with Restrictions 2‴, 3‴, and 4 imposed on the contemporaneous IRFs. The solid line shows the point-wise posterior medians, and the shaded bands show the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

E.8. Identification Using Local Projections

In this appendix, we test whether similar identification results for the supply chain disruption shock, as those shown in Figure 12, can be obtained using local projections (LPs). Unlike SVARs, LPs allow us to remain agnostic about the underlying multivariate system and to approximate the IRFs through a set of direct, linear regressions.

To approximate the impulse responses to a supply chain disruption shock, we apply the following regression model:

$$y_{i,t+k} = \beta_{i,k,0} \cdot ACR_t + \sum_{l=1}^{L} \boldsymbol{\beta}'_{i,k,l} \boldsymbol{y}_{t-l} + \boldsymbol{C}'_{i,k} \boldsymbol{\omega}_t + u_{i,k,t}, \tag{E.1}$$

where $1 \leq i \leq n$, $0 \leq k \leq K$, $\boldsymbol{y}_{t+k}$ is an $n \times 1$ vector of the k-step ahead forecasts of the same endogenous variables as in Section 4, $y_{i,t+k}$ is the value of the i-th variable in $\boldsymbol{y}_{t+k}$, $\boldsymbol{y}_{t-l}$ is an $n \times 1$ vector of the lagged endogenous variables, with the lag order denoted by l, $\boldsymbol{\omega}_t = [1, t]'$ is a 2×1 vector of a constant and a linear trend, and $u_{i,k,t}$ is the prediction error corresponding to the i-th variable. We are interested in the coefficients $\beta_{i,k,0}$ in Equation (E.1) for $k = 0, \ldots, K$, which encapsulate the dynamic responses of the i-th variable up to K months following a one standard deviation shock to the ACR index at the current time.

The model in Equation (E.1) can simply be estimated using the least squares method (Jordà, 2005), and the Newey-West standard errors (Newey and West, 1987) for $\beta_{i,k,0}$ can be used to quantify the uncertainty around the implied impulse responses. However, as the forecast horizon k increases, the estimates of $\beta_{i,k,0}$ become more susceptible to high variance, particularly given our small sample size. Hence, we apply the smooth local projections (SLP) method (Barnichon and Brownlees, 2019) by approximating the coefficient $\beta_{i,k,0}$ using a linear B-spline basis function expansion at horizon k, i.e.,

$$\beta_{i,k,0} \approx \sum_{m=1}^{M} b_{i,m} B_{i,m}(k), \tag{E.2}$$

where $B_{i,m}(k) : \mathbb{R} \to \mathbb{R}$ for $m = 1, \ldots, M$ represents a set of B-spline basis functions, and $b_{i,m}$ for $m = 1, \ldots, M$ denotes a set of scalar parameters. These parameters can be estimated using a penalized estimator that shrinks Equation (E.2) towards a quadratic polynomial at horizon k.

Following Barnichon and Brownlees (2019), we identify a supply chain disruption shock by controlling only for the lagged endogenous variables $\boldsymbol{y}_{t-l}$ (i.e., no contemporaneous values of other

endogenous variables are included as additional controls) and by estimating the regressions (E.1) from horizon $k = 0$. This approach is equivalent to placing the ACR index first in a recursive ordering, as all other endogenous variables can respond contemporaneously to shocks to the ACR index.[13] Additionally, we include four lags of the endogenous variables in the regressions (E.1) (i.e., $L = 4$), consistent with the maximum number of lags chosen for the SVAR model, as shown in Appendix E.3.

Figure E.46 shows the IRFs to a supply chain disruption shock. The black solid lines show the estimated impulse responses, while the gray-shaded areas represent the 90% confidence bands. In the short term, the responses of the endogenous variables are quantitatively similar to those shown in Figure 12; real GDP and product market tightness decline, spare capacity rises, and prices have not yet adjusted. Subsequently, as prices continue to rise, product market tightness starts to catch up, resulting in a continuous increase in real GDP accompanied by a simultaneous fall in spare capacity.[14] Compared to our SVAR estimation results, the responses of real GDP and product market tightness are not only statistically significant but also consistently elevated, evidenced by the lower boundary of the corresponding 90% confidence band remaining above the zero-response line for almost a year after the second-quarter mark. Additionally, the response of spare capacity is highly significant, with the upper boundary of the 90% confidence band staying well below the zero-response line when the estimated impulse response reaches its trough. Lastly, the ACR index remains elevated for more than six quarters post-impact.

[13] Our estimation results remain robust when we use a Cholesky ordered vector autoregression (VAR) model, with an ordering of the endogenous variables equal to *ACR*, *RealGDP*, *SpareCapacity*, *PCEGoodsPrice*, *ImportPrice*, and *ProductMarketTightness*. Our timing assumption posits that *ACR* has a contemporaneous effect on all variables ordered after it. In contrast, *RealGDP* only has a contemporaneous effect on *SpareCapacity* and the variables that follow, but not on *ACR*. Lastly, *SpareCapacity* affects both *ACR* and *RealGDP*, but only with a lag. Such a timing assumption is consistent with the zero restrictions we impose on the contemporaneous responses of the ACR index to aggregate demand and productive capacity shocks in Restrictions 1 and 2 when estimating the SVAR model. The VAR estimation results are available upon request.

[14] Benigno et al. (2022) apply the same SLP approach to estimate the response of U.S. consumer price inflation to shocks to the New York Fed's GSCPI. In identifying the structural shocks to the GSCPI, they control for the contemporaneous values of the year-on-year consumer price inflation rate and GDP growth rate, and estimate the LPs from horizon $k = 1$. This approach is equivalent to placing the GSCPI after consumer price inflation and GDP growth in a recursive ordering, implying that inflation does not respond contemporaneously to innovations in the GSCPI, while the GSCPI can respond to innovations in inflation contemporaneously. The estimated impulse response of U.S. consumer price inflation, as depicted in Figure 6 of their paper, is lower in magnitude compared to our estimate in Figure E.46. Similar comparison can also be found when we use the GSCPI as the measure indicating the state of the global supply chain in the SVAR model, as shown in Figure F.5.

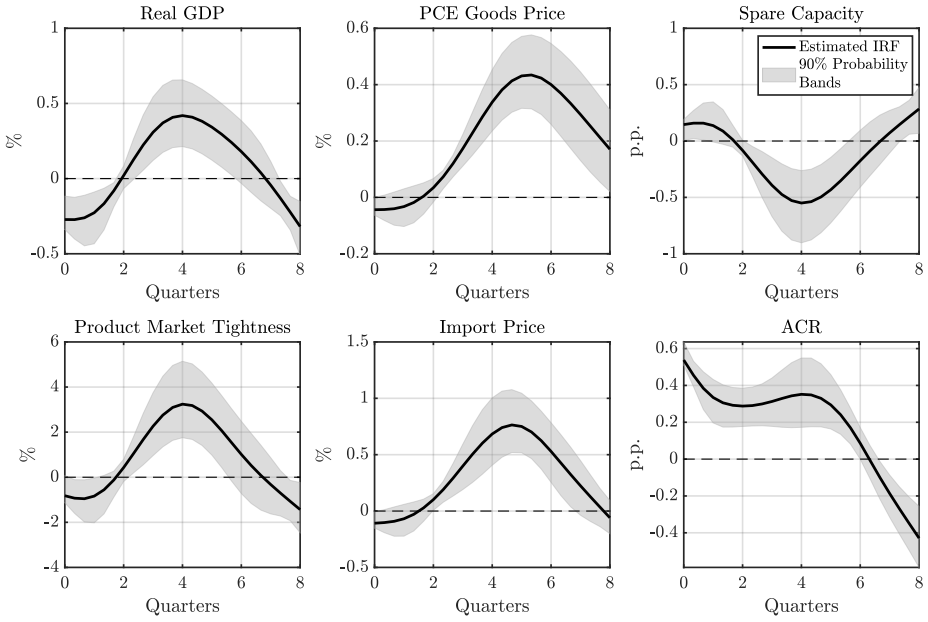

Figure E.46: IRFs to an Adverse Shock to the Supply Chain: Identification Through Local Projections

Notes. The IRFs to a one standard deviation supply chain disruption shock are computed using the SLP method, as described in Barnichon and Brownlees (2019). The shock is identified by controlling only for the lagged endogenous variables specified in Equation (E.1) and by estimating the regressions from horizon $k = 0$. This method is equivalent to a recursive identification scheme, with the ACR index positioned first in the ordering. The black solid lines show the estimated impulse responses, while the gray-shaded areas represent the 90% confidence bands.

F. Alternative Indices Measuring the State of the Global Supply Chain

In this appendix, we compare our ACR index to other popular indices that measure the state of the global supply chain in the existing literature: namely, the Harper Peterson Time Charter Rates Index (HARPEX), the Global Supply Chain Pressure Index (GSCPI) from the Federal Reserve Bank of New York (Benigno et al., 2022), and the Supply Disruptions Index (SDI) constructed by Smirnyagin and Tsyvinski (2022) and Liu et al. (2024). Our analysis reveals significant disparities between these indices, which affect the impact of supply chain disruptions on key macroeconomic indicators. In addition, we show that using the ACT index, as developed in Appendix B.3, in the causality assessment delivers quantitatively similar results to those obtained with the ACR index.

F.1. HARPEX

Shipping costs serve as a natural candidate to assess the health of the global supply chain (Benigno et al., 2022). They reflect the cost of transporting goods and can be influenced by various factors, including demand for shipping services, fuel prices, market speculation, and port efficiency. A sudden increase in shipping prices often indicates disruptions due to reduced shipping capacity or increased demand, signaling underlying issues in the flow of goods. Unlike shipping routes and itineraries, which respond to prevailing economic conditions with a significant lag, shipping costs react almost instantaneously to external shocks, whether due to a surge in demand, a shortage in productive capacity, or a disruption in the global supply chain.

Figures F.1, F.2, and F.3 present the estimation results with the HARPEX included in the SVAR model as the measure indicative of the state of the global supply chain. The HARPEX is a widely recognized composite indicator of container shipping rates in the time charter market across eight different classes of container ships (Benigno et al., 2022; Finck and Tillmann, 2022; Finck et al., 2024).[15] It is also integral to the construction of the New York Fed's GSCPI as a measure

[15] Finck and Tillmann (2022) and Finck et al. (2024) incorporate the HARPEX, along with the RWI/ISL container throughput index and the New York Fed's GSCPI, as endogenous variables in their SVAR models to reflect the state of international container shipping and global supply chains. Their identification of a global supply chain shock employs conventional sign restrictions on the impulse responses, augmented by narrative information related to the Tōhoku earthquake in 2011, the Suez Canal blockage in 2021, and the Shanghai backlog in 2022.

of cross-border transportation costs. In estimating the SVAR, we removed the zero restrictions while retaining the sign restrictions to discipline the IRFs, consistent with the observation that shipping prices respond to demand and capacity shocks more flexibly and quickly than our ACR index.[16]

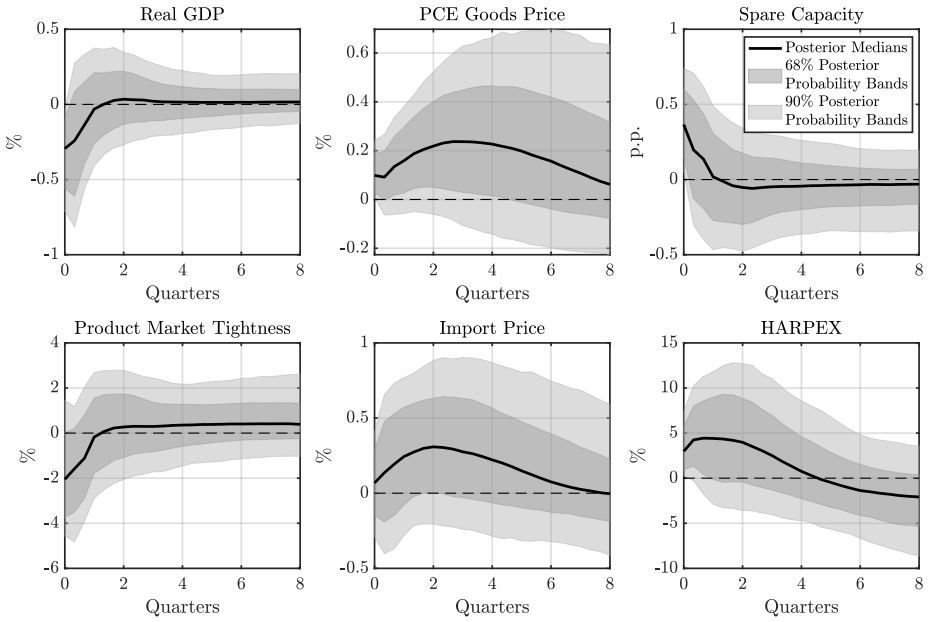

Figure F.1: IRFs to an Adverse Shock to the Supply Chain: The HARPEX

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using the HARPEX and Restrictions 1′, 2′, and 3′. The solid line shows the point-wise posterior medians, and the shaded bands represent the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

By examining the IRFs to an adverse supply chain disturbance, we find that, similar to the results obtained with the ACR index, the shock leads to stagflation accompanied by an increase in spare capacity. However, the confidence bands for the PCE goods and import prices are noticeably wider than those shown in Figure 12 of the main text. Furthermore, supply chain disturbances no longer account for the largest fraction of unexpected fluctuations in PCE goods and import prices in both the medium and long term; instead, shocks to productive capacity play a major role in driving these price fluctuations. Regarding the historical decomposition of

In contrast, we include our ACR index as the variable indicative of the state of the global supply chain, and our identification of a global supply chain disturbance involves sign restrictions derived from our theoretical model, as described in Section 3 of the main text.

[16]We choose not to impose any sign restrictions on the impulse responses of the HARPEX to demand and capacity shocks, since our theoretical model does not provide such restrictions, and transportation costs are idiosyncratic. Alternatively, we could augment the current model with a full-fledged transportation sector, where interactions between producers, shipowners, and retailers determine the transportation cost, hence providing theoretical predictions on its responses to demand and capacity shocks. We leave this for future research.

U.S. goods inflation, the HARPEX yields results quantitatively similar to those obtained in the baseline estimation. Notably, supply chain disruptions were the main driver behind both the rising inflation in 2021 and the subsequent fall in inflation in the second half of 2022.

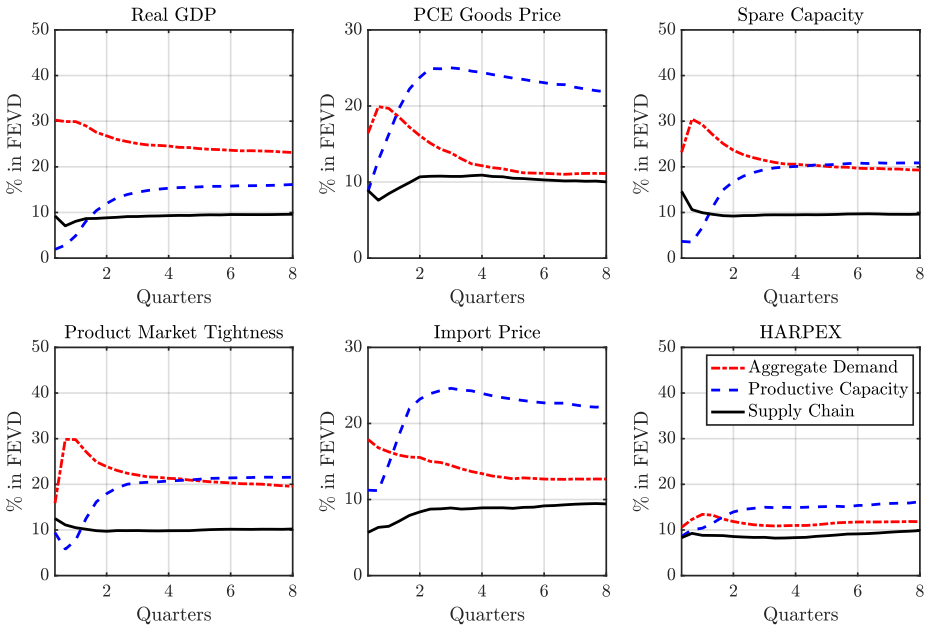

Figure F.2: FEVD from the SVAR: The HARPEX

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated using the HARPEX and Restrictions 1′, 2′, and 3′, and based on 100,000 independent importance sampling draws.

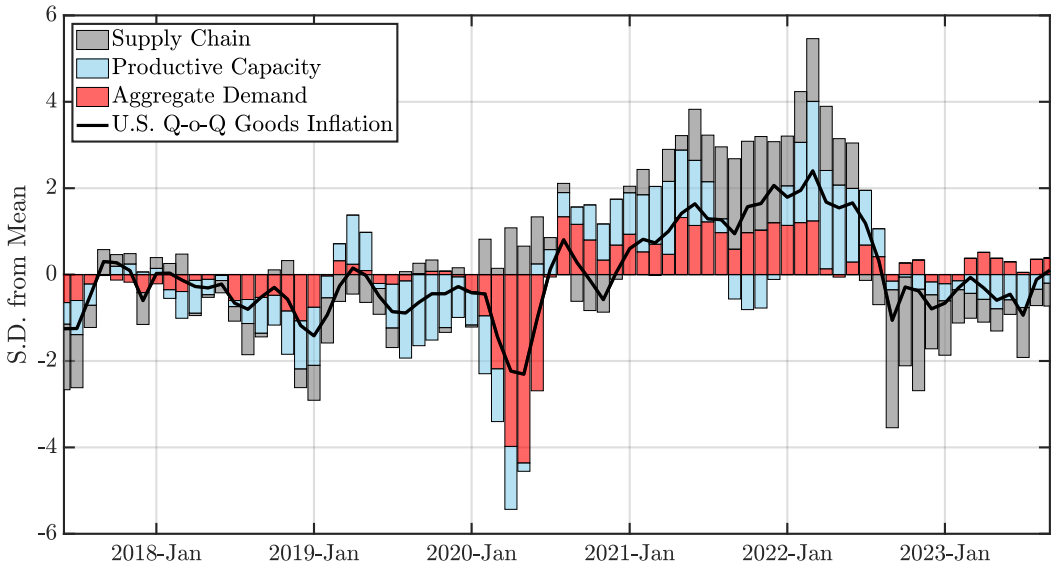

Figure F.3: HD of U.S. Quarter-on-Quarter Goods Inflation: The HARPEX

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with the HARPEX included as the measure representing the state of the global supply chain, and with Restrictions 1′, 2′, and 3′ imposed on the contemporaneous IRFs. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

F.2. GSCPI

Next, we compare our ACR index with the New York Fed's GSCPI to highlight the differences in the assessments of the state of the global supply chain. The GSCPI uses information on cross-border transportation costs and the sub-components of the country-specific manufacturing PMI (e.g., "delivery times") to infer supply chain disruptions.[17] As discussed in the main text, the GSCPI provides a reading on disruptions to the supply chain inferred from: (i) transportation costs that are subject to variations from factors unrelated to supply chain disruptions, (ii) information from purchasing managers that might reflect subjective views or speculative predictions rather than actual disturbances, and (iii) the PMI, which cannot distinguish whether increased delivery times result from supply chain disruptions or issues within the production process. Our ACR index is derived using accurate satellite data, closely tracking port congestion worldwide – a manifestation of global supply chain disruptions (Transportation Research Board Executive Committee, 2006). This index of global port congestion provides an alternative, precise measure to gauge the state of the global supply chain.

Figure F.4 plots the ACR and the GSCPI indices from January 2017 to September 2023. Prior to 2020, the dynamics of these two indices were aligned. However, at the onset of the COVID-19 pandemic in early 2020, the GSCPI saw a substantial increase and remained high in the first half of the year. From late 2020 onward, both series exhibited a parallel rise until January 2022. In an influential paper, di Giovanni et al. (2022) link the early 2020 surge in the GSCPI to the beginning of the lockdown in China and the subsequent decline in the second half of 2020 to the partial reopening of China and Europe.

Our port congestion index suggests that the initial lockdown in China did not cause congestion to a degree that would lead to significant global supply chain disruptions. Similarly, the reopening of China and Europe did not notably alleviate port congestion. Thus, the fluctuations in the GSCPI are likely influenced by abrupt shifts in demand and market expectation, as well as managerial misperceptions of supply chain issues, as indicated by PMI surveys, not a surprise given the high levels of uncertainty about COVID-19 during the first weeks of the pandemic. Notably, the two series diverged again in early 2022; the ACR index remained high, while the

[17] IHS Markit calculates the suppliers' delivery times using the survey responses in the PMI. Specifically, participating purchasing managers are asked if it takes their suppliers more or less time to provide inputs to their factories on average. The percentages of companies reporting an improvement, deterioration, or no change in delivery times are then weighted to derive the index.

GSCPI started to decrease. We argue that the elevated levels of the ACR index during the first half of 2022 were largely due to the stringent containment measures in China, which remained in effect during this period and exerted significant pressure on the global supply chain.

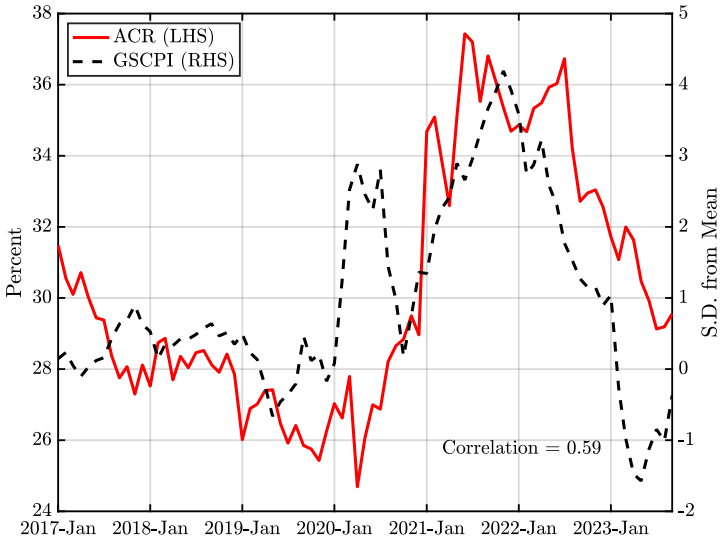

Figure F.4: ACR vs. GSCPI

Notes. Figure F.4 plots the ACR index (solid red line) against the GSCPI (dashed black line) for the sample period from January 2017 to September 2023. The ACR index is computed using the AIS data of container ships and our IMA-DBSCAN algorithm, as detailed in Appendix B. The GSCPI is retrieved from the website of the Federal Reserve Bank of New York. The ACR index is measured in percentage terms, while the GSCPI is in standard deviations from the mean. Both series have been seasonally adjusted.

Now, we examine the implications of the discrepancies between the ACR and GSCPI indices for the inferences drawn by the SVAR model regarding the causal effects of supply chain disruptions. Figure F.5 plots the IRFs to an adverse supply chain disturbance, with the GSCPI incorporated in the SVAR model as the measure indicative of the state of the global supply chain. To retain consistency, we remove the zero restrictions imposed on the contemporaneous responses of the GSCPI to demand and productive capacity shocks, while retaining the sign restrictions to discipline the IRFs. While the median responses of the endogenous variables are similar in shape to the baseline responses obtained using the ACR index, the confidence bands are noticeably wider, suggesting less precision in the estimates. In particular, the median responses of the PCE goods and import prices are estimated to be lower in magnitude, and the 68% posterior probability bands for these prices also encompass the zero-response line. In contrast, the baseline estimates, as depicted in Figure 12, consistently show these responses well above the zero-response line.

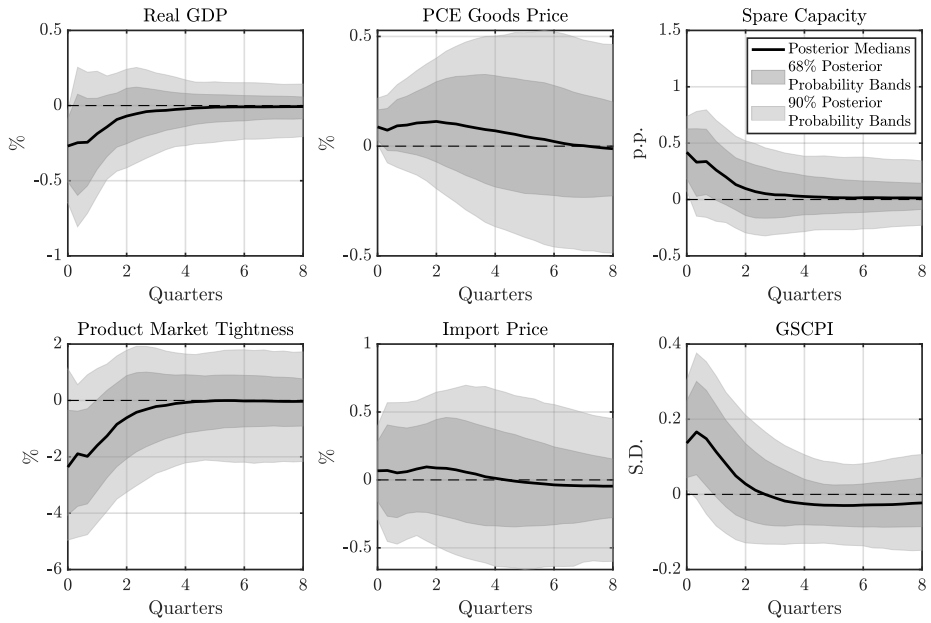

Figure F.5: IRFs to an Adverse Shock to the Supply Chain: The GSCPI

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using the GSCPI and Restrictions 1′, 2′, and 3′. The solid line shows the point-wise posterior medians and the shaded bands represent the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

Figure F.6 presents the proportion of forecast error variance explained by each of the three structural shocks. Consistent with the decomposition observed using the ACR index (as depicted in Figure 13), shocks to aggregate demand continue to be the primary source of unexpected fluctuations in real GDP, spare capacity, and product market tightness, even when the GSCPI is included in the estimation. However, there is a notable difference in the impact of supply chain shocks. While these shocks account for a considerable portion of the unforeseen variations in PCE goods and import prices over extended periods in analyses using the ACR index, their influence is markedly less pronounced when the GSCPI is employed.

Finally, Figure F.7 shows the cumulative historical contribution of each of the three identified shocks to U.S. goods inflation when the GSCPI is incorporated into the estimation. Similar to the results based on the ACR index, those using the GSCPI attribute the sudden decline in inflation at the onset of the COVID-19 pandemic largely to adverse demand shocks, while simultaneously noting that supply chain disruptions exert upward pressure on inflation. The difference emerges in late 2020, when the estimation using the GSCPI attributes the decline in inflation to improvements in the global supply chain, consistent with the behavior of the GSCPI

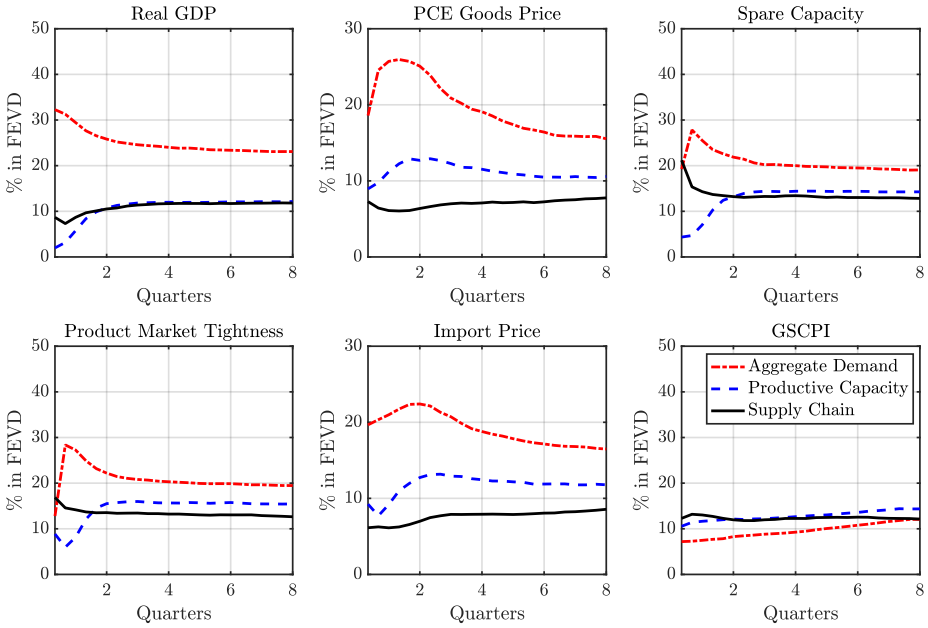

Figure F.6: FEVD from the SVAR: The GSCPI

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated using the GSCPI and Restrictions 1′, 2′, and 3′, and based on 100,000 independent importance sampling draws.

in the second half of 2020, as observed in Figure F.4. Meanwhile, the sustained increase in inflation from late 2020 onward is explained by a combination of the three shocks, with supply chain shocks playing a dominant role in driving inflation dynamics during late 2021 to early 2022.

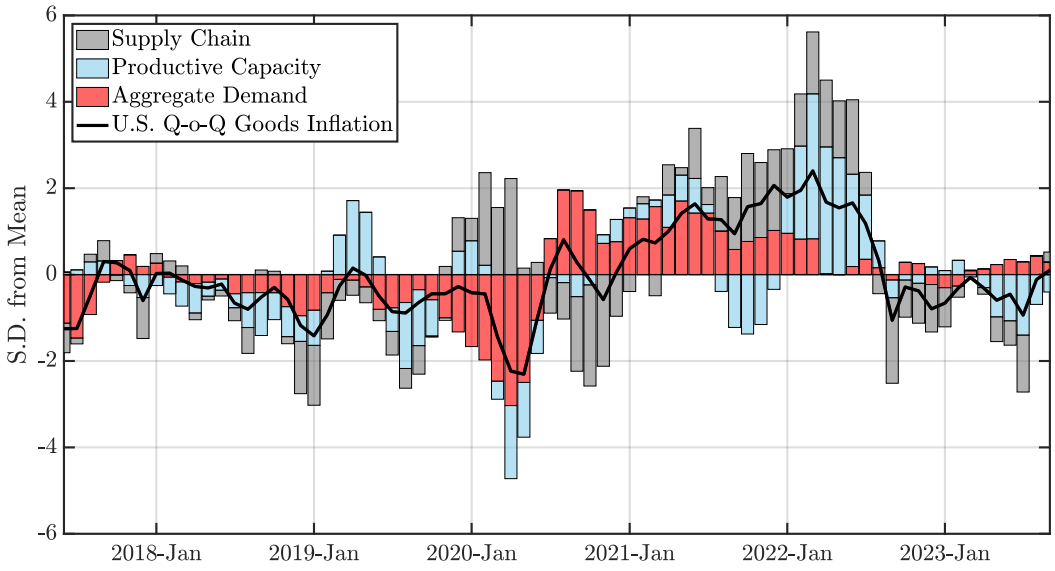

Figure F.7: HD of U.S. Quarter-on-Quarter Goods Inflation: The GSCPI

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with the GSCPI included as the measure representing the state of the global supply chain, and with Restrictions 1′, 2′, and 3′ imposed on the contemporaneous IRFs. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

F.3. SDI

In addition to the HARPEX and GSCPI, several other indices utilize either more advanced techniques (e.g., machine learning) or more granular data (e.g., bills of lading). For instance, Smirnyagin and Tsyvinski (2022) and Liu et al. (2024) leverage the S&P Global Panjiva dataset, a comprehensive repository of U.S. seaborne import records, to derive the U.S. Supply Disruptions Index (SDI). They identify supply chain disruptions by observing regular and active consignee-shipper relationships over quarterly periods; a disruption is marked when a consistently active relationship becomes inactive for a quarter before resuming.

Although this identification strategy zeroes in on disruptions within established trading relationships, it might misclassify intentional decisions to halt trading for reasons unrelated to supply chain disruptions. For example, a consignee might temporarily cease orders due to diminished demand, while a shipper might temporarily halt supply due to a lack of labor rather than a genuine supply chain disruption. Parsing out these scenarios based solely on the activity of consignee-shipper relationships can be intricate. Furthermore, while the SDI provides invaluable insights into U.S. imports and excels in generating asset pricing predictions, our ACR index uses port congestion to infer supply chain disruptions, accurately reflecting changes in the regular flow of goods worldwide.

Figure F.8 plots both the ACR and the SDI indices. The SDI remained low before 2020, experienced a significant but brief spike in 2020, followed by a moderate and sustained increase throughout 2021, and then began to decline in 2022. The estimation results, derived by incorporating the SDI in the SVAR model as the measure indicative of the state of the global supply chain and imposing only sign restrictions to discipline the IRFs, are depicted in Figures F.9, F.10, and F.11. Echoing the findings with the GSCPI, the responses of the PCE goods and import prices to a supply chain shock, as estimated with the SDI, are less precise. Additionally, only a small fraction of the unexpected fluctuations in these prices is attributed to supply chain disturbances. The historical decomposition of U.S. goods inflation indicates that the initial drop in inflation at the onset of the pandemic was primarily due to a collapse in demand. Improvements in the global supply chain contributed to lowering inflation in the second half of 2020. Furthermore, a combination of demand and supply shocks drove the subsequent increase in inflation from late 2020 onward, with supply chain disruptions playing a major role during late 2021 to early 2022.

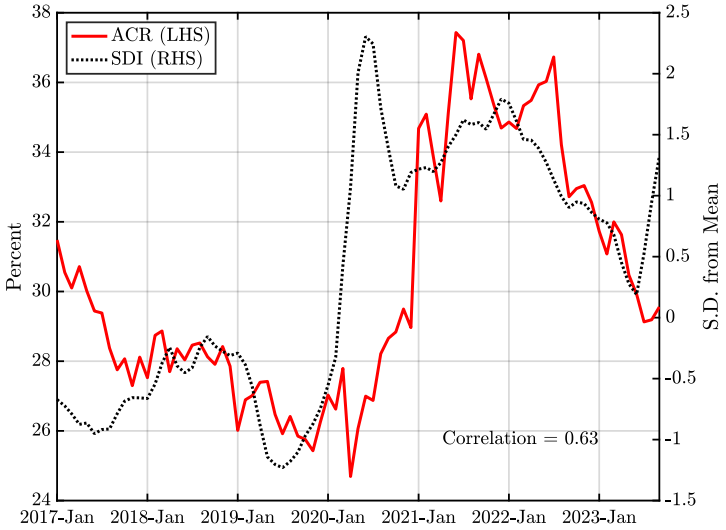

Figure F.8: ACR vs. SDI

Notes. Figure F.8 plots the ACR index (solid red line) against the SDI (black dotted line) for the sample period from January 2017 to September 2023. The ACR index is computed using the AIS data of container ships and our IMA-DBSCAN algorithm, as detailed in Appendix B. The SDI is retrieved from the author's website. The ACR index is measured in percentage terms, while the SDI is in standard deviations from the mean. Both series have been seasonally adjusted.

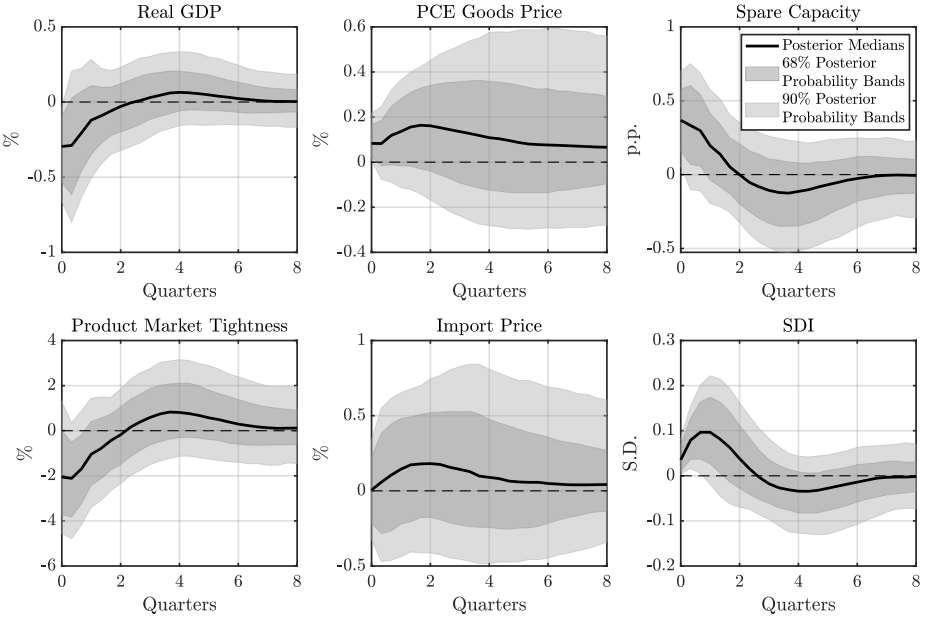

Figure F.9: IRFs to an Adverse Shock to the Supply Chain: The SDI

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using the SDI and Restrictions 1′, 2′, and 3′. The solid line shows the point-wise posterior medians, and the shaded bands represent the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

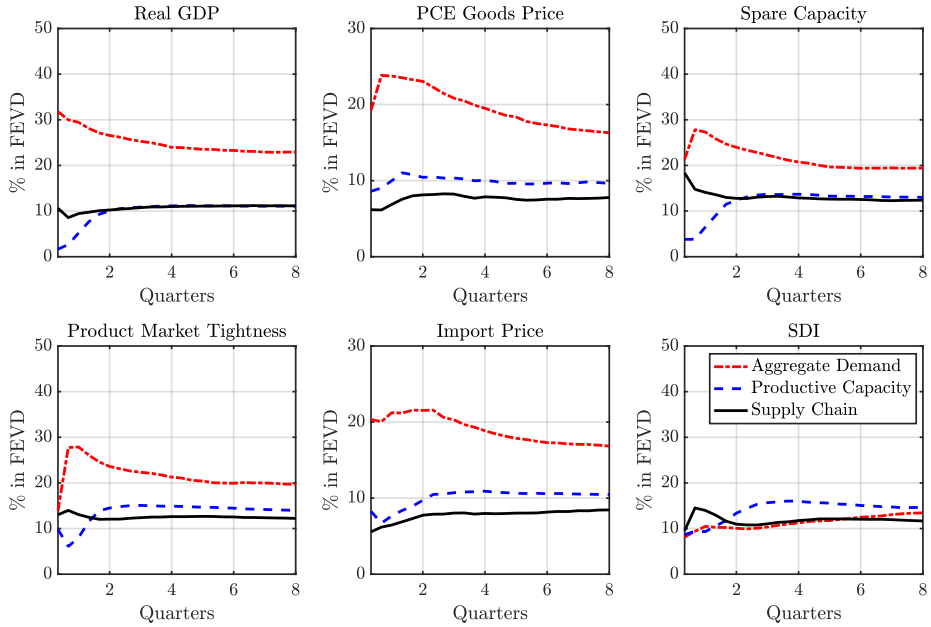

Figure F.10: FEVD from the SVAR: The SDI

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated using the SDI and Restrictions 1′, 2′, and 3′, and based on 100,000 independent importance sampling draws.

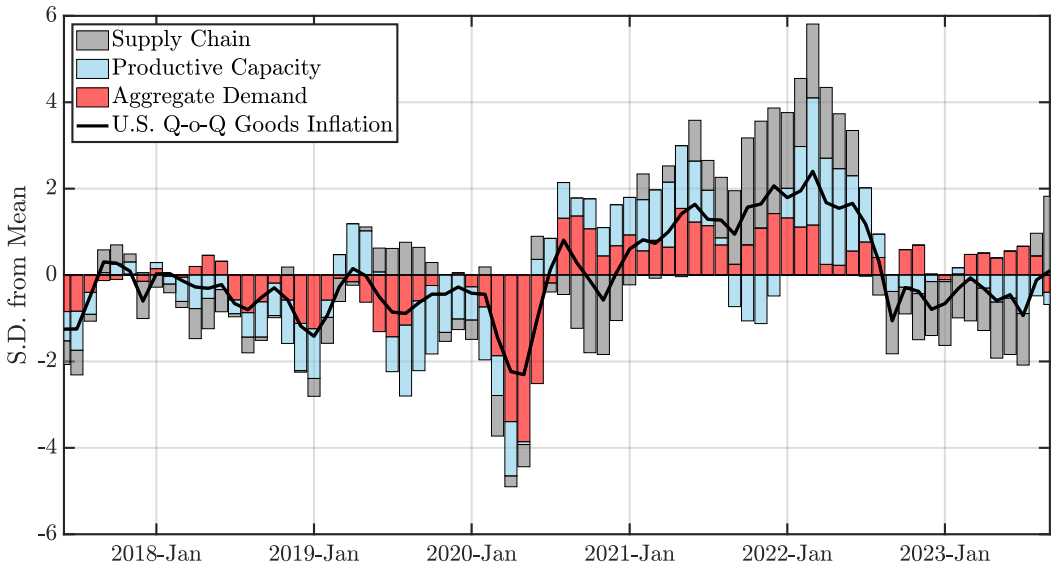

Figure F.11: HD of U.S. Quarter-on-Quarter Goods Inflation: The SDI

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with the SDI included as the measure representing the state of the global supply chain, and with Restrictions 1′, 2′, and 3′ imposed on the contemporaneous IRFs. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

F.4. Other Indices in the Literature

In addition to the aforementioned indices, Kilian et al. (2023) introduce a new monthly index of the volume of container trade to and from North America. While this index is useful for identifying shocks to domestic demand and foreign demand for U.S. manufactured goods, it is not suitable for our purposes in identifying global supply chain disturbances.

Another notable measure is the text-based index of supply disruptions developed by Burriel et al. (2023) using newspaper data, following the methodology of Baker et al. (2016). While this index focuses exclusively on supply-side events, it is not immune to measurement errors inherent in word definitions. For instance, a disruption in "supply" is distinct from one in the "supply chain," as the former may result from factors such as a labor supply shortage.

The Kiel Trade Indicator (Stamer, 2021), which utilizes the same AIS data as our ACR index, offers an alternative perspective on the global supply chain strain by estimating imports and exports and tracking the movement of container ships at major ports, as well as freight on stationary ships. However, this indicator faces several challenges. First, the calculation of TEUs is problematic because the draft of a container ship does not reliably indicate its loading status, given that loading and unloading operations can occur simultaneously. Second, variations in imports and exports are influenced by several factors, including changes in demand and market expectations, and are therefore subject to issues similar to those of indices based on transportation costs. Third, the methodology for estimating cargo capacity tied up at ports does not distinguish between the mooring positions of container ships (i.e., berth or anchorage), potentially leading to an overestimation of port congestion, as argued in Talley (2009), Talley and Ng (2016), and Brancaccio et al. (2024).

Such an overestimation of port congestion also appears in the calculation of container ship dwell times at ports, as discussed by Fuchs and Wong (2022). Specifically, using information on the ship's speed and navigation status (whether it is under way using its engine or moored at a pier), they define a ship's dwell time as the period it spends moored at a pier with zero speed. As they explain in their paper, this measure is conservative because it does not account for: (i) the time a ship spends navigating within the port area as it prepares to moor at a pier, and (ii) the potential waiting time outside the port area at an anchorage before it navigates to the port.

F.5. ACR for the Major Ports Along the Trans-Pacific Route

In addition to the aggregate ACR index, we also develop a targeted ACR index for the major container ports along the Trans-Pacific route, which is arguably the most important and busiest shipping route connecting East Asia (mostly China) and the U.S., accounting for nearly 30 million TEUs of cargo transported across the Pacific Ocean every year. We use this index in the SVAR estimation as a robustness check, given that congestion in ports along the Trans-Pacific route might exert a disproportionately larger impact on the U.S. economy than other major ports around the world. Consequently, the shock identified using this targeted index might represent a supply chain disruption more specific to the U.S. than the disturbance identified through the aggregate ACR index.

Figure F.12 plots this targeted ACR index against the aggregate ACR index at a monthly frequency. Despite the overall trends of the two congestion indices aligning closely with each other during the sample period, there are more and larger fluctuations in the targeted ACR index, especially during the pandemic period.

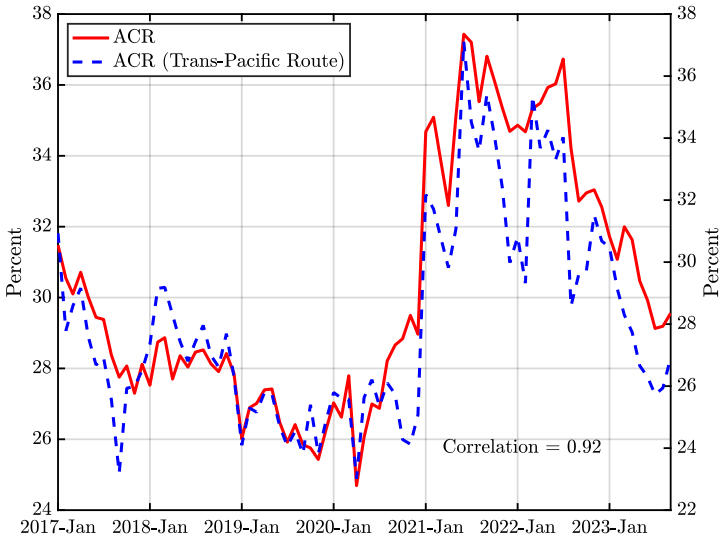

Figure F.12: Aggregate ACR vs. Targeted ACR

Notes. Figure F.12 plots the aggregate ACR index (solid red line) against a targeted ACR index for the major ports along the Trans-Pacific route (dashed blue line) during the sample period from January 2017 to September 2023. Both indices are computed using the AIS data of container ships and our IMA-DBSCAN algorithm, as detailed in Appendix B. While the aggregate ACR index accounts for congestion at the top 50 global container ports, the targeted ACR index aggregates congestion at the Port of Busan in South Korea, the Ports of Hanshin and Keihin in Japan, the Ports of Los Angeles and Long Beach in the U.S., and the Ports of Dalian, Dongguan, Guangzhou Harbor, Hong Kong, Kaohsiung, Ningbo-Zhoushan, Qingdao, Shenzhen, Shanghai, Tianjin, Xiamen, and Yang Shan in China. Both indices are measured in percentage terms and have been seasonally adjusted.

Figure F.13 shows the impulse responses of the selected macro aggregates to a supply chain disturbance, identified off Restrictions 1, 2, and 3, with the targeted ACR index included in the SVAR model as the measure indicative of the health of the supply chain specific to the U.S. Similar to the findings with the aggregate ACR index, the shock leads to stagflation alongside a rise in spare capacity. Furthermore, Figure F.14 shows the proportion of forecast error variance explained by each of the three structural shocks. Unlike those presented in Figure 13, supply chain shocks no longer explain the largest fraction of unexpected variations in PCE goods and import prices in the medium and long term; instead, capacity shocks become the predominant driving force of these price fluctuations. Lastly, the historical decomposition of U.S. goods inflation, as shown in Figure F.15, is almost identical to the one using the aggregate ACR index, implying that our narratives regarding the different drivers of U.S. goods inflation during 2020-2023 remain consistent even when we use a measure of port congestion specific to the U.S.

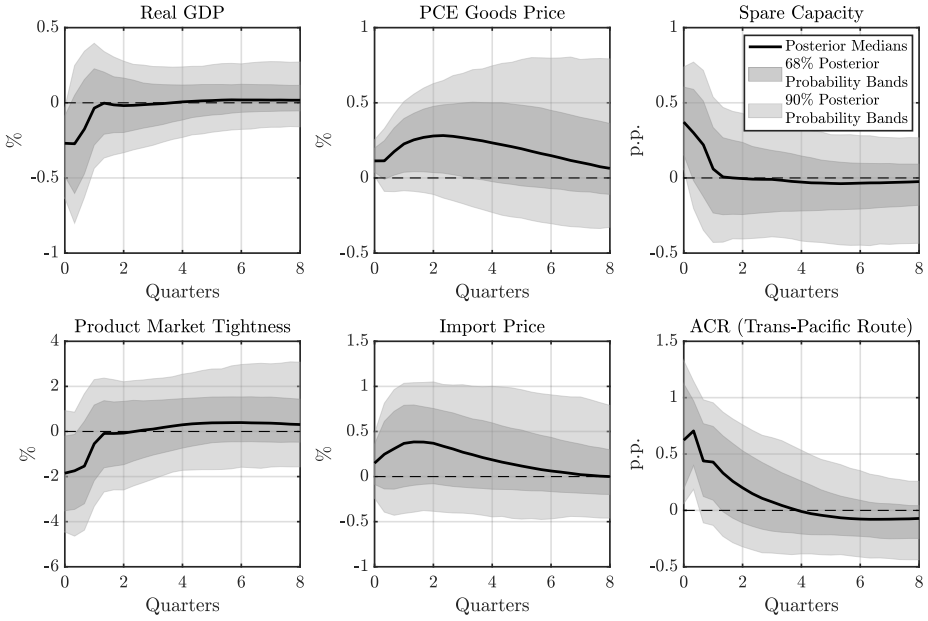

Figure F.13: IRFs to an Adverse Shock to the Supply Chain: The Targeted ACR Index

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using a targeted ACR index for the major container ports along the Trans-Pacific route and Restrictions 1, 2, and 3. The solid line shows the point-wise posterior medians, and the shaded bands represent the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

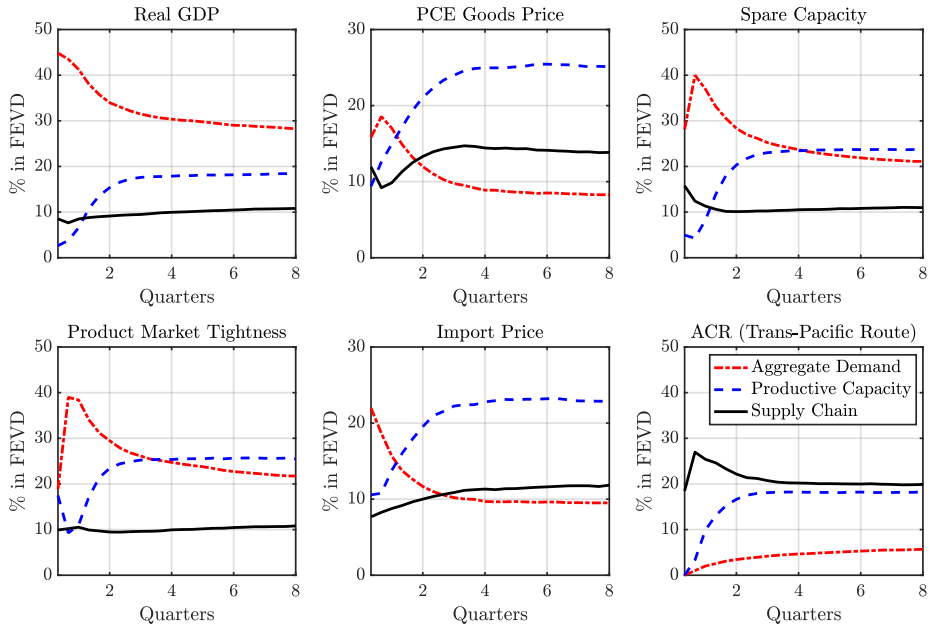

Figure F.14: FEVD from the SVAR: The Targeted ACR Index

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated using a targeted ACR index for the major container ports along the Trans-Pacific route and Restrictions 1, 2, and 3, and based on 100,000 independent importance sampling draws.

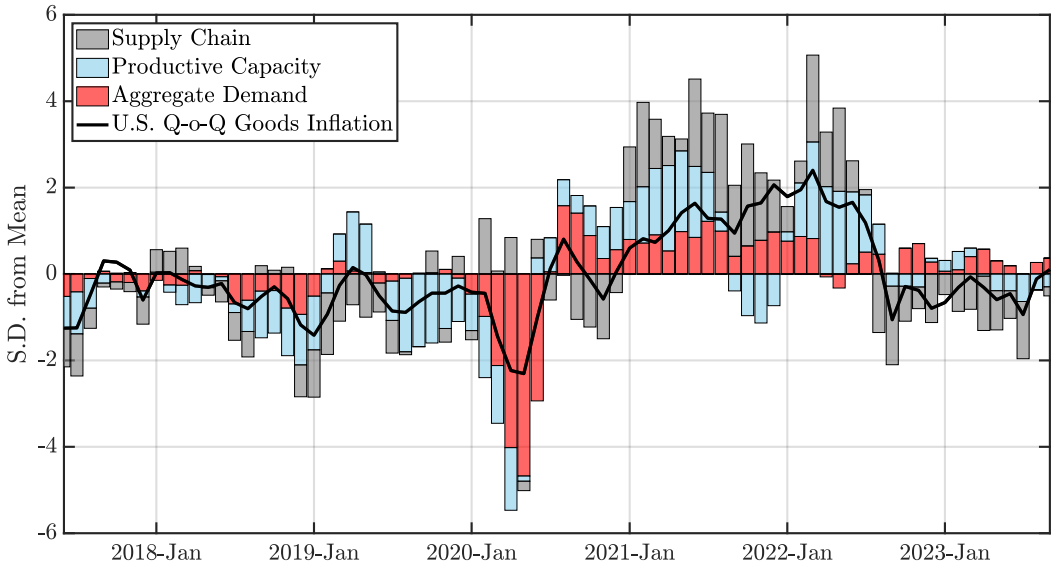

Figure F.15: HD of U.S. Quarter-on-Quarter Goods Inflation: The Targeted ACR Index

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with a targeted ACR index for the major container ports along the Trans-Pacific route included as the measure indicative of the state of the supply chain specific to the U.S., and with Restrictions 1, 2, and 3 imposed on the contemporaneous IRFs. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

F.6. ACT

As introduced in Appendix B.3, we define an alternative congestion metric for ports, namely the Average Congestion Time (ACT). This metric measures the average number of hours a container ship waits in an anchorage area of a port before docking at a berth, with the measurement weighted by the relative number of ship visits to each port. Figure F.16 plots the ACT index alongside the ACR index at a monthly frequency.

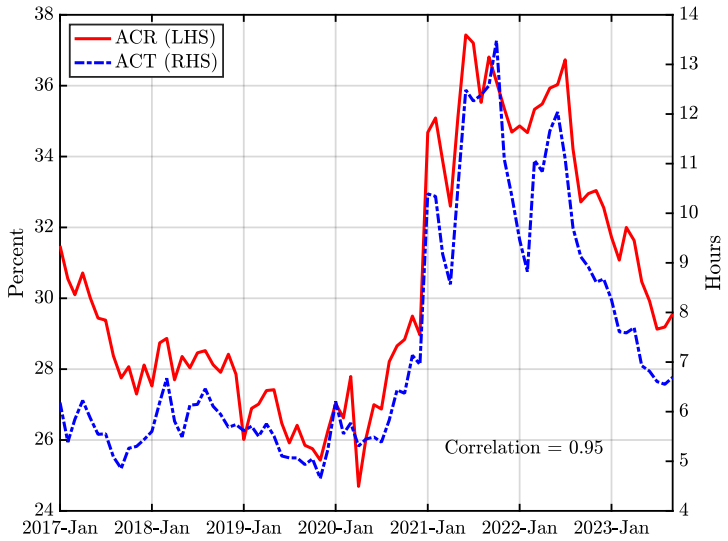

Figure F.16: ACR vs. ACT

Notes. Figure F.16 plots the ACR index (solid red line) against the ACT index (dashed-dotted blue line) for the sample period from January 2017 to September 2023. Both the ACR and ACT indices are computed using the AIS data of container ships and our IMA-DBSCAN algorithm, as detailed in Appendix B. The ACR index is measured in percentage terms, while the ACT index is in hours. Both series have been seasonally adjusted.

As shown in Figure F.16, the two congestion indices closely co-move with each other, exhibiting a correlation of 0.95. Not surprisingly, as illustrated in Figures F.17, F.18, and F.19, using the ACT index in the causality assessment delivers results quantitatively similar to those obtained with the ACR index.

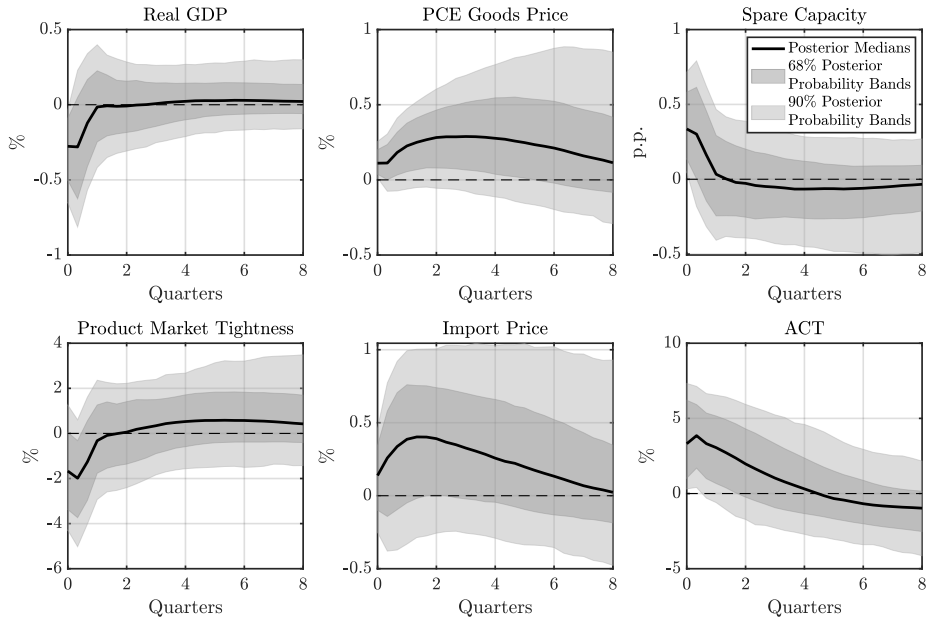

Figure F.17: IRFs to an Adverse Shock to the Supply Chain: The ACT Index

Notes. The IRFs to a one standard deviation adverse shock to the supply chain are identified using the ACT index and Restrictions 1, 2, and 3. The solid line shows the point-wise posterior medians, and the shaded bands represent the 68% and 90% equal-tailed point-wise posterior probability bands. The figure is based on 100,000 independent importance sampling draws.

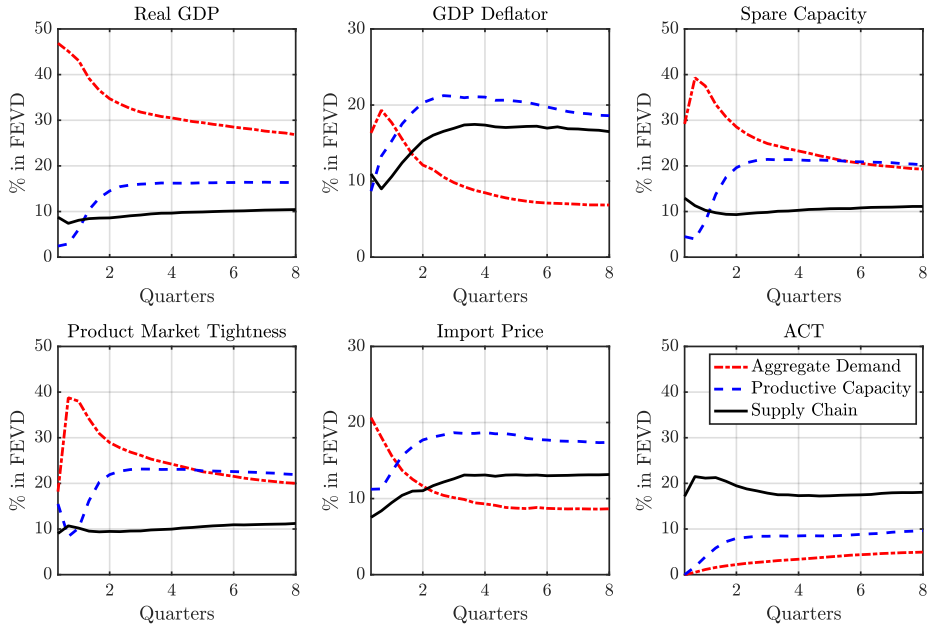

Figure F.18: FEVD from the SVAR: The ACT Index

Notes. Each line presents the median fraction of the forecast error variance for each endogenous variable, explained by each of the three identified structural shocks at various time horizons. The FEVD is estimated using the ACT index and Restrictions 1, 2, and 3, and based on 100,000 independent importance sampling draws.

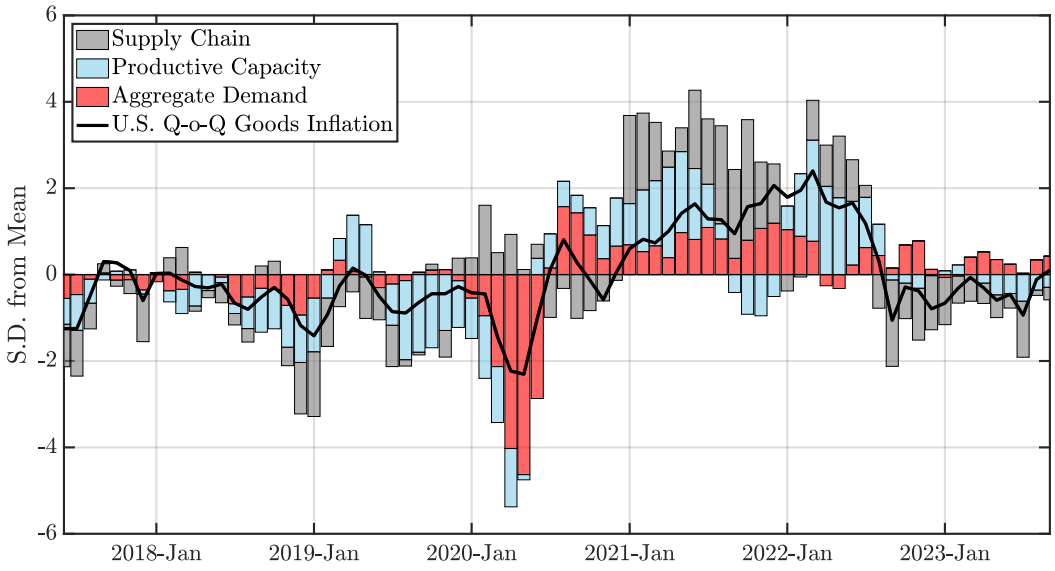

Figure F.19: HD of U.S. Quarter-on-Quarter Goods Inflation: The ACT Index

Notes. The solid line represents the standardized goods inflation rate in the U.S., i.e., the quarter-on-quarter growth of the PCE goods price index. The shaded bars represent the corresponding standardized cumulative historical contribution of shocks to aggregate demand, productive capacity, and the supply chain to goods inflation. The shocks are identified using the SVAR specification in Equation (28), with the ACT index included as the measure representing the state of the global supply chain, and with Restrictions 1, 2, and 3 imposed on the contemporaneous IRFs. The figure is derived from the posterior medians, based on 100,000 independent importance sampling draws.

G. Priors and Identification in the TVAR

G.1. Priors

Our formulation of the prior in the TVAR model follows Bańbura et al. (2010), Mumtaz and Zanetti (2012), and Pizzinelli et al. (2020), and the same prior has been applied to the parameters in both the supply chain disrupted ($\mathbb{D}$) and undisrupted ($\mathbb{U}$) regimes. Specifically, we write the TVAR model in Equation (33) compactly as a system of multivariate regressions:

$$\boldsymbol{y} = (\boldsymbol{M}_{\mathbb{D}}\boldsymbol{x}_{\mathbb{D}} + \boldsymbol{u}_{\mathbb{D}})\boldsymbol{I} + (\boldsymbol{M}_{\mathbb{U}}\boldsymbol{x}_{\mathbb{U}} + \boldsymbol{u}_{\mathbb{U}})(\boldsymbol{1}_{T\times T} - \boldsymbol{I}), \tag{G.1}$$

where $\boldsymbol{y} = [\boldsymbol{y}_1 \ \dots \ \boldsymbol{y}_T]$ is an $n \times T$ matrix, $\boldsymbol{x}_{\mathbb{D}} = [\boldsymbol{x}_{\mathbb{D},1} \ \dots \ \boldsymbol{x}_{\mathbb{D},T}]$ is an $m \times T$ matrix with $\boldsymbol{x}_{\mathbb{D},t} = [\boldsymbol{y}'_{t-1} \ \dots \ \boldsymbol{y}'_{t-L} \ \boldsymbol{\omega}'_t]'$, $\boldsymbol{x}_{\mathbb{U}} = [\boldsymbol{x}_{\mathbb{U},1} \ \dots \ \boldsymbol{x}_{\mathbb{U},T}]$ is an $m \times T$ matrix with $\boldsymbol{x}_{\mathbb{U},t} = [\boldsymbol{y}'_{t-1} \ \dots \ \boldsymbol{y}'_{t-L} \ \boldsymbol{\omega}'_t]'$, $\boldsymbol{\omega}_t = [1, t]'$ is a 2×1 vector of a constant and a linear trend, $\boldsymbol{u}_{\mathbb{D}} = [\boldsymbol{\Sigma}_{\mathbb{D}}^{1/2}\boldsymbol{\varepsilon}_1 \ \dots \ \boldsymbol{\Sigma}_{\mathbb{D}}^{1/2}\boldsymbol{\varepsilon}_T]$ is an $n\times T$ matrix, $\boldsymbol{u}_{\mathbb{U}} = [\boldsymbol{\Sigma}_{\mathbb{U}}^{1/2}\boldsymbol{\varepsilon}_1 \ \dots \ \boldsymbol{\Sigma}_{\mathbb{U}}^{1/2}\boldsymbol{\varepsilon}_T]$ is an $n \times T$ matrix, $\boldsymbol{\Sigma}_{\mathbb{D}}$ and $\boldsymbol{\Sigma}_{\mathbb{U}}$ are the variance-covariance matrices, $\boldsymbol{I} = diag[I_1 \ \dots \ I_T]$ is a $T \times T$ diagonal matrix, $\boldsymbol{M}_{\mathbb{D}} = [\boldsymbol{B}'_{\mathbb{D},1} \ \dots \ \boldsymbol{B}'_{\mathbb{D},L} \ \boldsymbol{C}'_{\mathbb{D}}]$ and $\boldsymbol{M}_{\mathbb{U}} = [\boldsymbol{B}'_{\mathbb{U},1} \ \dots \ \boldsymbol{B}'_{\mathbb{U},L} \ \boldsymbol{C}'_{\mathbb{U}}]$ are two $n \times m$ matrices containing the TVAR coefficients associated with each regime, and $m = nL + 2$. Given Equation (G.1), for each regime $r \in \{\mathbb{D}, \mathbb{U}\}$, we assume that the prior distribution of the parameter vector, $vec(\boldsymbol{M}_r)$, has a Normal-Inverse-Wishart (NIW) conjugate form.[18] Such a form can be written as:

$$\begin{aligned} vec(\boldsymbol{M}_r)|\boldsymbol{\Sigma}_r &\sim N\left(vec(\boldsymbol{M}_r^0), \boldsymbol{\Sigma}_r \otimes \boldsymbol{\Omega}_r^0\right), \\ \boldsymbol{\Sigma}_r &\sim IW\left(\boldsymbol{S}_r^0, \alpha_r^0\right), \end{aligned} \tag{G.2}$$

where $vec(\boldsymbol{M}_r^0)$ is the prior mean of the parameter vector, $\boldsymbol{\Omega}_r^0$ controls the tightness around this prior, $\boldsymbol{S}_r^0$ is the prior scale matrix of the Inverse-Wishart (IW) distribution, and α_r^0 denotes the prior degrees of freedom. Essentially, the prior in Equation (G.2) is a generalization of the Minnesota prior discussed in Litterman (1986) and assumes that the endogenous variables follow a random walk or an AR(1) process. This is based on the idea that recent lags provide more reliable information on the dynamics of the system and therefore the estimation should assign them a higher weighting. Unlike the original formulation in Litterman (1986), however, the prior in Equation (G.2) does not assume a diagonal, fixed, and known variance-covariance matrix,

[18] $vec(\cdot)$ denotes the operator that stacks the columns of a matrix into a vector.

making it more suitable for our structural analysis.

The Normal-Inverse-Wishart prior implies that, while the prior expectations and variances of the coefficient matrices for the constant and linear trend, $\boldsymbol{C}_r$, are diffuse, those associated with the autoregressive matrices, $\boldsymbol{B}_{r,l}$, can be written as follows:

$$\begin{aligned}\mathbb{E}[(\boldsymbol{B}_{r,l})_{i,j}] &= \begin{cases}\beta^0_{r,i}, & \text{if } i=j,\ l=1;\\ 0, & \text{otherwise;}\end{cases}\\ \mathbb{V}[(\boldsymbol{B}_{r,l})_{i,j}] &= \lambda\sigma_i^2/\sigma_j^2,\end{aligned} \tag{G.3}$$

where $\beta^0_{r,1},\ldots,\beta^0_{r,n}$ are the prior means of the autoregressive coefficients, $\sigma_1,\ldots,\sigma_n$ are the prior error standard deviations, and the hyper-parameter λ controls the overall tightness of the prior distribution such that a larger λ corresponds to a looser prior. As described in Bańbura et al. (2010) and commonly used in the literature of Bayesian SVARs, the prior moments in Equation (G.3) can be implemented by adding $T_{r,d}$ dummy observations $\boldsymbol{y}_{r,d}$ and $\boldsymbol{x}_{r,d}$ to the system of regressions in Equation (G.1) that correspond to each regime, with $\boldsymbol{y}_{r,d}$ and $\boldsymbol{x}_{r,d}$ satisfying the following structures:

$$\boldsymbol{y}_{r,d} = \begin{bmatrix} diag[\beta^0_{r,1}\sigma_1 \ \ldots \ \beta^0_{r,n}\sigma_n]/\lambda \\ \mathbf{0}_{n(L-1)\times n} \\ diag[\sigma_1 \ \ldots \ \sigma_n] \\ \mathbf{0}_{2\times n}\end{bmatrix},\ \boldsymbol{x}_{r,d} = \begin{bmatrix} J_L \otimes diag[\sigma_1 \ \ldots \ \sigma_n]/\lambda & \mathbf{0}_{nL\times 1} & \mathbf{0}_{nL\times 1} \\ \mathbf{0}_{n\times nL} & \mathbf{0}_{n\times 1} & \mathbf{0}_{n\times 1} \\ \mathbf{0}_{1\times nL} & \xi & 0 \\ \mathbf{0}_{1\times nL} & 0 & \xi\end{bmatrix},$$

where $J_L = diag[1 \ \ldots \ L]$ and the hyper-parameter ξ controls the prior on the constant and the linear trend such that a small number makes the prior uninformative. Subsequently, the prior moments in Equation (G.2) are simply functions of $\boldsymbol{y}_{r,d}$ and $\boldsymbol{x}_{r,d}$, which are given by:

$$\begin{aligned}\boldsymbol{M}^0_r &= \boldsymbol{y}_{r,d}\boldsymbol{x}'_{r,d}(\boldsymbol{x}_{r,d}\boldsymbol{x}'_{r,d})^{-1},\\ \boldsymbol{\Omega}^0_r &= (\boldsymbol{x}_{r,d}\boldsymbol{x}'_{r,d})^{-1},\\ \boldsymbol{S}^0_r &= (\boldsymbol{y}_{r,d}-\boldsymbol{M}^0_r\boldsymbol{x}_{r,d})(\boldsymbol{y}_{r,d}-\boldsymbol{M}^0_r\boldsymbol{x}_{r,d})',\\ \alpha^0_r &= T_{r,d}-m.\end{aligned}$$

With the Normal-Inverse-Wishart prior being conjugate, the conditional posterior distribution of the parameter vector is also Normal-Inverse-Wishart (Bańbura et al., 2010; Mumtaz and

Zanetti, 2012):

$$vec(\boldsymbol{M}_r)|\boldsymbol{\Sigma}_r, \boldsymbol{y} \sim N\left(vec(\tilde{\boldsymbol{M}}_r), \boldsymbol{\Sigma}_r \otimes (\tilde{\boldsymbol{x}}_r \tilde{\boldsymbol{x}}_r')^{-1}\right),$$
$$\boldsymbol{\Sigma}_r|\boldsymbol{y} \sim IW\left(\tilde{\boldsymbol{S}}_r, T_{r,d} + 2 + T - m\right),$$

where the parameters associated with the posterior are given by:

$$\tilde{\boldsymbol{M}}_r = \tilde{\boldsymbol{y}}_r \tilde{\boldsymbol{x}}_r' (\tilde{\boldsymbol{x}}_r \tilde{\boldsymbol{x}}_r')^{-1},$$
$$\tilde{\boldsymbol{S}}_r^0 = (\tilde{\boldsymbol{y}}_r - \tilde{\boldsymbol{M}}_r \tilde{\boldsymbol{x}}_r)(\tilde{\boldsymbol{y}}_r - \tilde{\boldsymbol{M}}_r \tilde{\boldsymbol{x}}_r)',$$

in which the terms $\tilde{\boldsymbol{y}}_r$ and $\tilde{\boldsymbol{x}}_r$ are the matrices of $\boldsymbol{y}_r$ and $\boldsymbol{x}_r$ augmented with the dummy observations $\boldsymbol{y}_{r,d}$ and $\boldsymbol{x}_{r,d}$ respectively.[19]

Following Mumtaz and Zanetti (2012) and Pizzinelli et al. (2020), we obtain the values of the prior mean of each autoregressive coefficient, $\beta_{r,i}^0$, as well as the prior error standard deviation, σ_i, from the OLS estimation of a univariate AR(1) model for each endogenous variable. In addition, we set $\lambda = 0.25$ to ensure fast lag decay toward zero. Finally, in terms of the prior distribution of $\overline{ACR}$, we assume that it is normally distributed, with the mean set at the median of the ACR series and the standard deviation calibrated to deliver a Markov Chain Monte Carlo acceptance rate of approximately 70% to 75%.

G.2. Identification Using the PFA

Following Uhlig (2005) and Mountford and Uhlig (2009), the identification scheme we employ in the study of the state-dependent effects of a contractionary monetary policy shock amounts to finding an impulse vector a that minimizes a given criterion function $f(\cdot)$ on the space of all impulse vectors. This function penalizes positive impulse responses of real GDP, PCE goods price, product market tightness, and import price as well as negative impulse responses of the federal funds rate and spare capacity at horizons $k = 1, \ldots, K$, while satisfying the zero restriction imposed on the impulse response of the ACR index at horizon $k = 1$.[20] The scheme is applied separately for the observations in each regime. Hence, for simplicity, we drop the regime-specific notation $r \in \{\mathbb{D}, \mathbb{U}\}$ in the following description.

[19] $\boldsymbol{y}_r$ is the part of $\boldsymbol{y}$ that is associated with regime $r \in \{\mathbb{D}, \mathbb{U}\}$.

[20] Our identification results using the PFA are also robust to removing the zero restriction imposed on the impulse response of the ACR index at horizon $k = 1$.

The PFA is implemented numerically as follows. Define the penalty function as:

$$f(x) = \begin{cases} x, & \text{if } x \leq 0; \\ 100x, & \text{if } x > 0, \end{cases} \tag{G.4}$$

which penalizes positive responses in linear proportion and rewards negative responses in linear proportion, albeit at a slope one hundred times smaller than the slope for penalties on the positive side. For the true VAR coefficients, let $r_{j,a}(k),\ k = 1, \ldots, K$ be the impulse response of variable j and let σ_j be the standard deviation of the series for variable j. Let $\iota_j = -1$ if j is the index of the federal funds rate or spare capacity in the data vector, and $\iota_j = 1$ if j is the index of real GDP, PCE goods price, product market tightness, or import price in the data vector. Define the contractionary monetary policy impulse vector as that impulse vector a, which minimizes the total penalty $\varphi(a)$ subject to the zero restriction imposed on the impulse response of the ACR index at horizon $k = 1$:

$$\varphi(a) = \sum_{j \in \left\{ \begin{array}{c} \text{“federal funds rate,”} \\ \text{“real GDP,”} \\ \text{“PCE goods price,”} \\ \text{“spare capacity,”} \\ \text{“product market tightness,”} \\ \text{“import price”} \end{array} \right\}} \left[\sum_{k=1}^{K} f\left(\iota_j \frac{r_{j,a}(k)}{\sigma_j} \right) \right].$$

The re-scaling by σ_j is necessary to make the deviations across different impulse responses comparable to each other. Notice that the sign of the penalty direction is flipped for the federal funds rate and spare capacity. Since the true VAR is unknown, we find the contractionary monetary policy vector for each draw from the posterior. Such a step involves numerical minimization, and we keep all the draws and accordingly calculate all the corresponding impulse vectors. As a result, the IRFs in the main text are calculated based on these.

G.3. Posterior and Identified Regimes

The posterior distribution of the threshold $\overline{ACR}$ is plotted in Figure G.1, while the time series of the identified regimes using the median of such a posterior is plotted in Figure G.2.

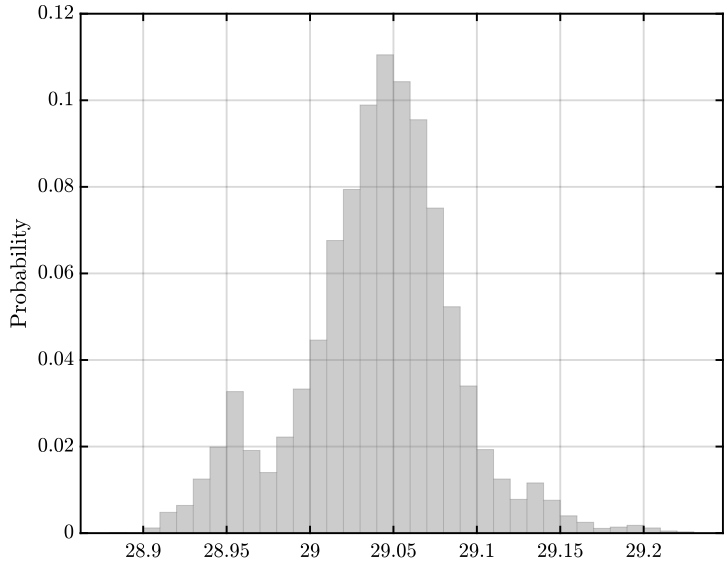

Figure G.1: Posterior Distribution of the Threshold $\overline{ACR}$

Notes. The figure plots the posterior distribution of the ACR threshold value, i.e., $\overline{ACR}$, based on 10,000 independent draws.

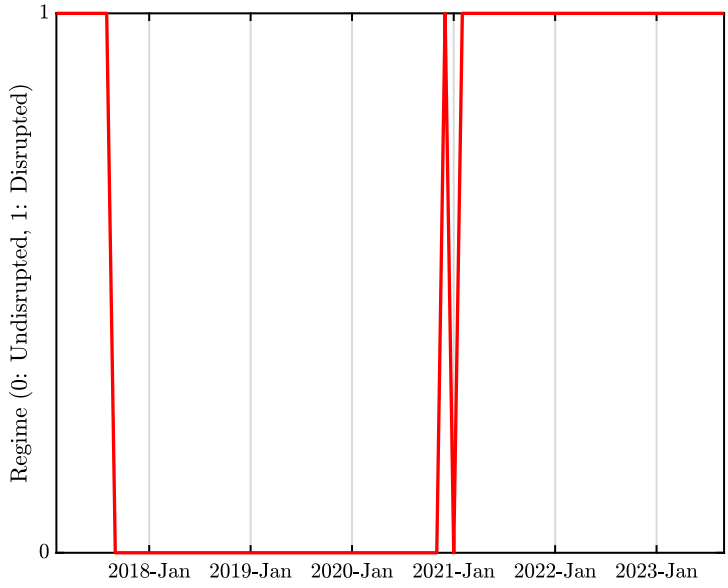

Figure G.2: Regimes Based on the Median of the Posterior $\overline{ACR}$

Notes. The solid red line, switching from zero to one, represents the current regime as identified by the median of the posterior distribution of the ACR threshold, i.e., $median(\overline{ACR}) = 29.04\%$. The value of one corresponds to the supply chain disrupted ($\mathbb{D}$) regime, while the value of zero corresponds to the supply chain undisrupted ($\mathbb{U}$) regime.

H. Robustness of TVAR Results

In this appendix, we conduct several robustness checks of our identification results presented in Section 5 regarding the state-dependent effects of monetary policy. In Figure H.1, we first show that our state-dependence results are robust to using the Wu-Xia shadow federal funds rate (Wu and Xia, 2016). This series provides an estimated federal funds rate during periods when the zero lower bound constrains it and, therefore, reflecting the effects of unconventional monetary policy.

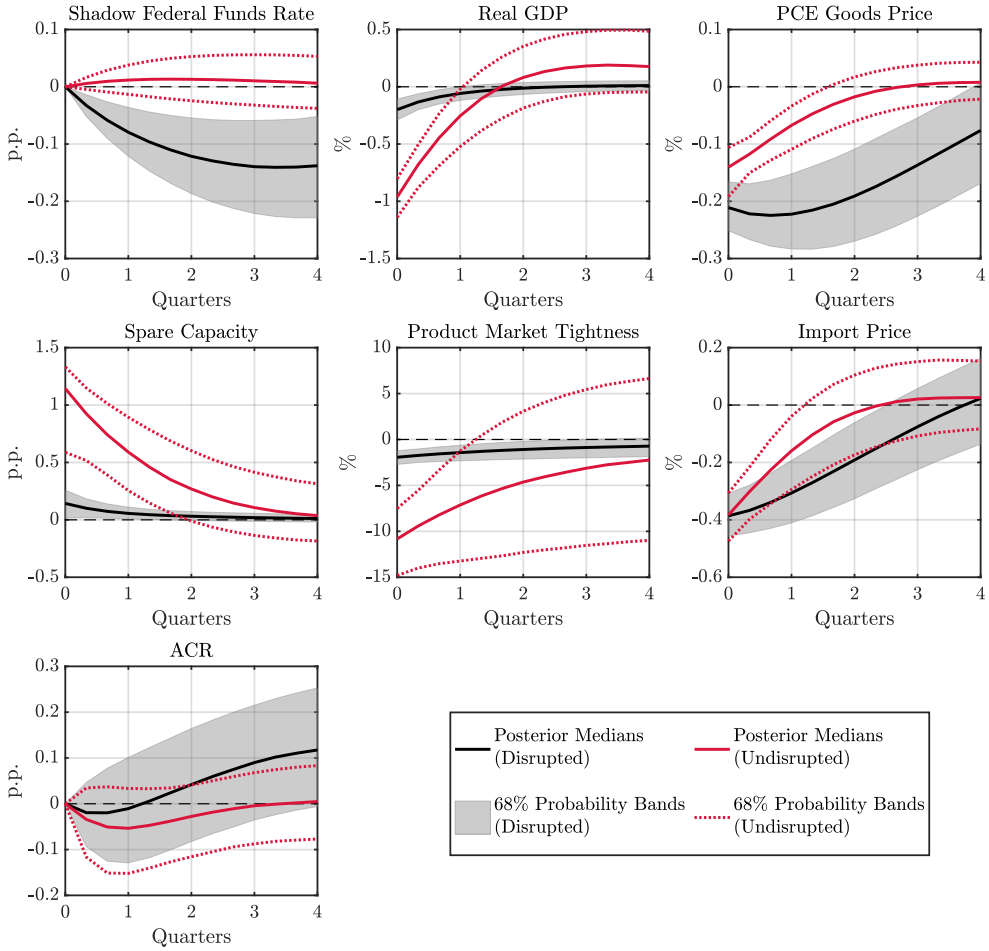

Figure H.1: State-Dependent Effects of a Contractionary Monetary Policy Shock: Wu-Xia Shadow Federal Funds Rate

Notes. The figure shows the IRFs to a one standard deviation contractionary monetary policy shock identified using a TVAR specification in Equation (33), with the Wu-Xia shadow federal funds rate (Wu and Xia, 2016) included to reflect the stance of U.S. monetary policy, as well as Restriction 4 imposed on the IRFs of the endogenous variables, for both the supply chain disrupted and undisrupted regimes. The solid black (solid red) line shows the point-wise posterior medians, and the shaded black area (dotted red lines) depicts the 68% equal-tailed point-wise posterior probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 10,000 independent draws from the posterior.

In Figure H.2, we show that our state-dependence results are robust to dropping the zero restriction imposed on the on-impact response of the ACR index to the contractionary monetary policy shock.

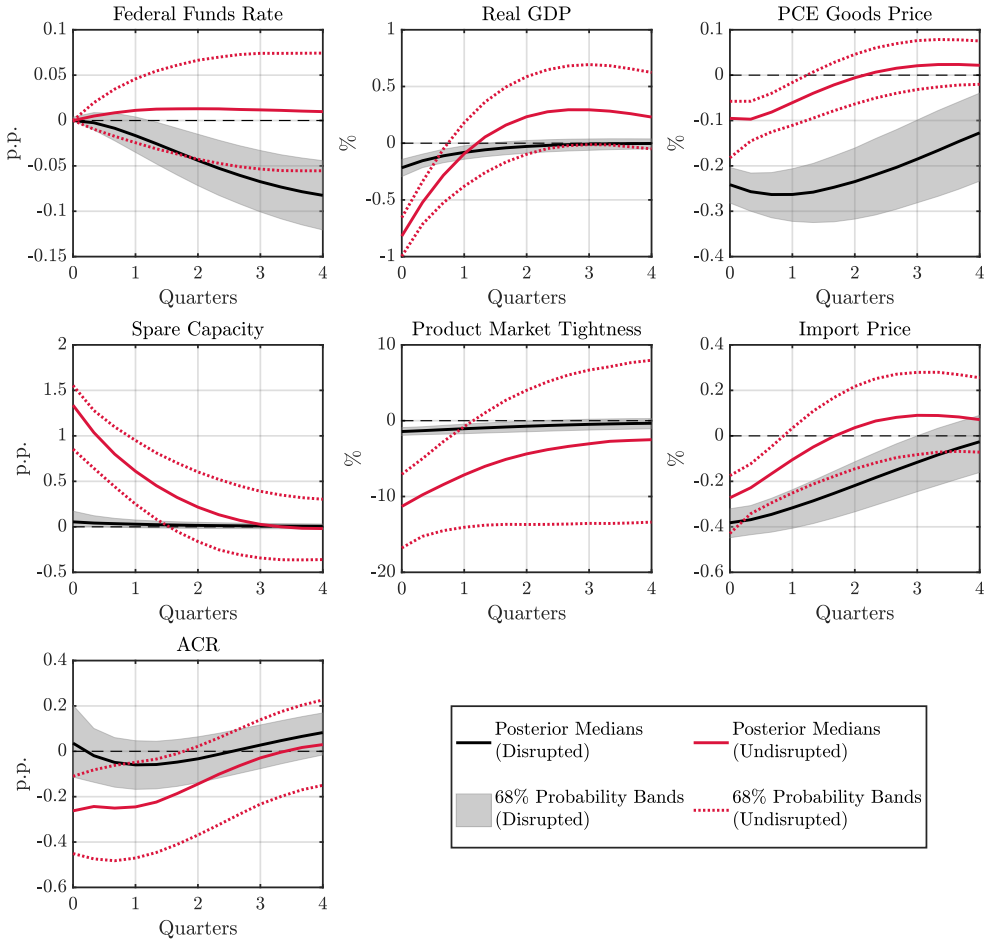

Figure H.2: State-Dependent Effects of a Contractionary Monetary Policy Shock: No Zero Restriction on the ACR Index

Notes. The figure shows the IRFs to a one standard deviation contractionary monetary policy shock, identified using a TVAR specification in Equation (33), but without the zero restriction imposed on the on-impact response of the ACR index, for both the supply chain disrupted and undisrupted regimes. The solid black (solid red) line shows the point-wise posterior medians, and the shaded black area (dotted red lines) depicts the 68% equal-tailed point-wise posterior probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 10,000 independent draws from the posterior.

Also, in Figures H.3 and H.4, we show that our state-dependence results are robust to considering different lag structures, i.e., two or three lags. We do not consider four lags or beyond due to parameter uncertainty resulting from our limited sample length. We also show in Figure H.5 that the results are robust when a looser prior is undertaken in the estimation, i.e., $\lambda = 0.5$.

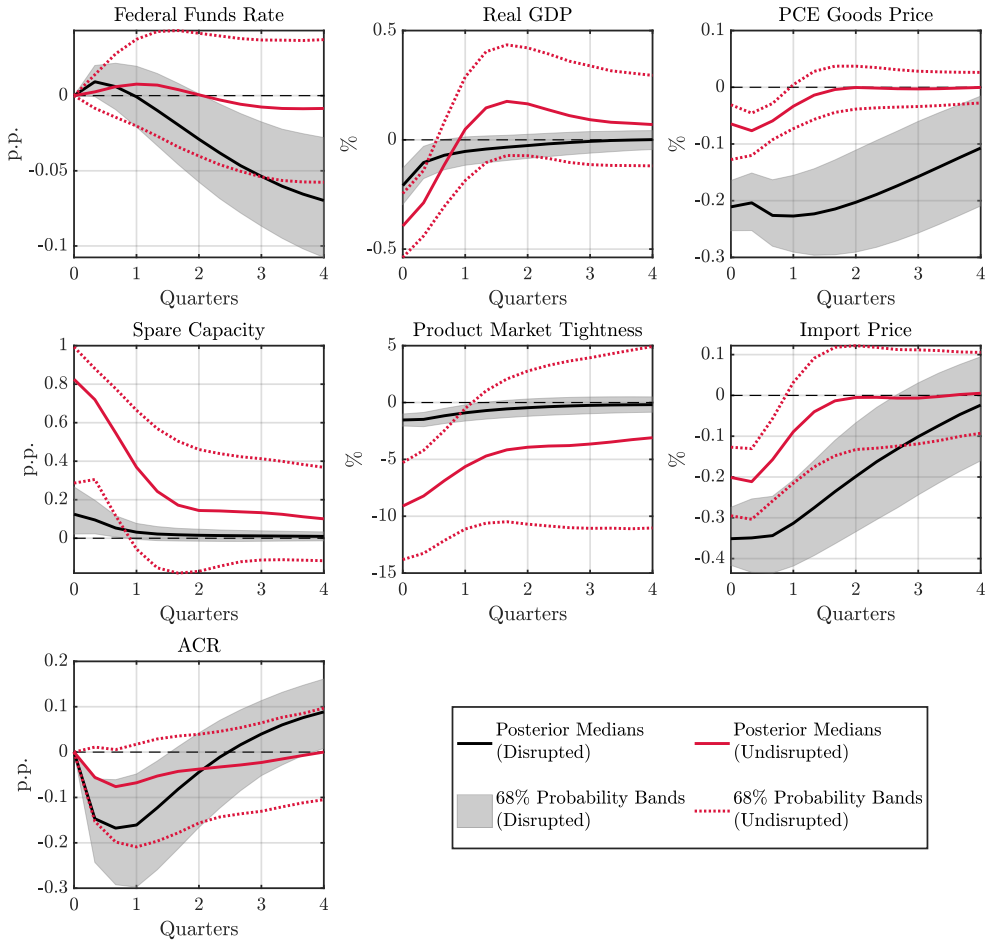

Figure H.3: State-Dependent Effects of a Contractionary Monetary Policy Shock: Two Lags

Notes. The figure shows the IRFs to a one standard deviation contractionary monetary policy shock identified using a TVAR specification in Equation (33) with two lags, as well as Restriction 4 imposed on the IRFs of the endogenous variables, for both the supply chain disrupted and undisrupted regimes. The solid black (solid red) line shows the point-wise posterior medians, and the shaded black area (dotted red lines) depicts the 68% equal-tailed point-wise posterior probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 10,000 independent draws from the posterior.

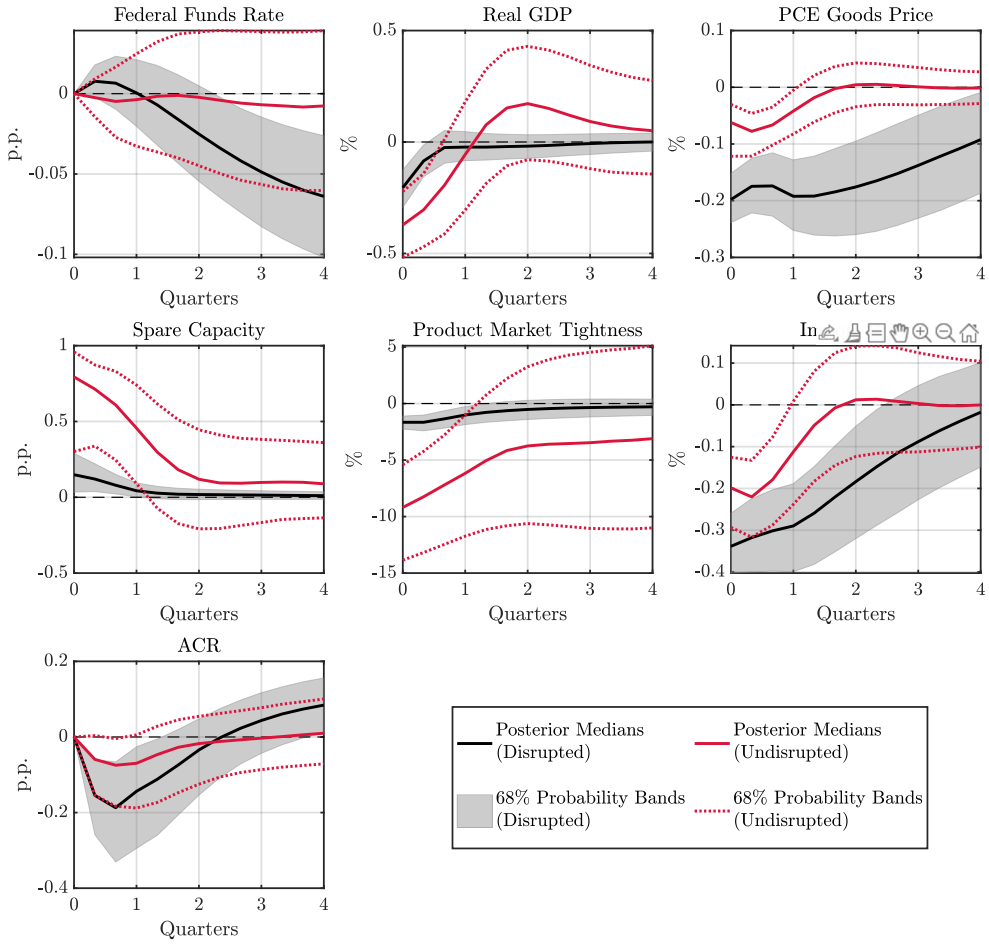

Figure H.4: State-Dependent Effects of a Contractionary Monetary Policy Shock: Three Lags

Notes. The figure shows the IRFs to a one standard deviation contractionary monetary policy shock identified using a TVAR specification in Equation (33) with three lags, as well as Restriction 4 imposed on the IRFs of the endogenous variables, for both the supply chain disrupted and undisrupted regimes. The solid black (solid red) line shows the point-wise posterior medians, and the shaded black area (dotted red lines) depicts the 68% equal-tailed point-wise posterior probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 10,000 independent draws from the posterior.

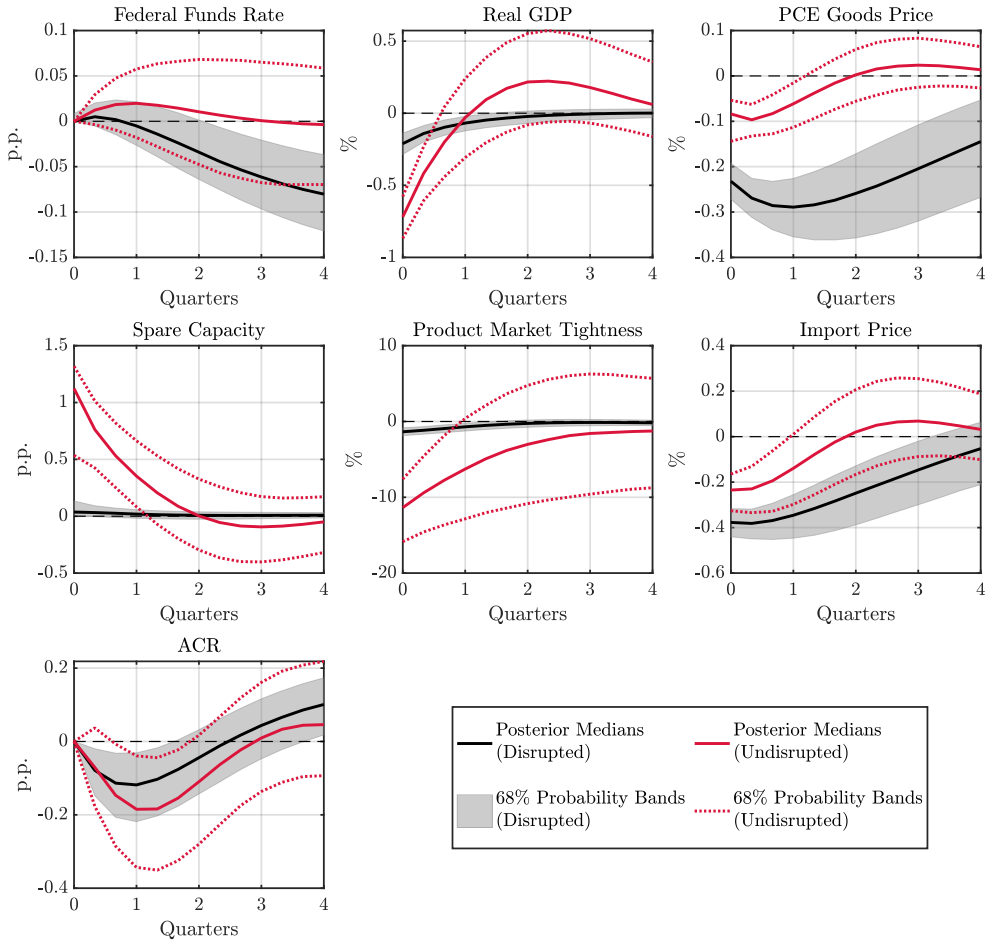

Figure H.5: State-Dependent Effects of a Contractionary Monetary Policy Shock: Looser Prior

Notes. The figure shows the IRFs to a one standard deviation contractionary monetary policy shock identified using a TVAR specification in Equation (33) with $\lambda = 0.5$, as well as Restriction 4 imposed on the IRFs of the endogenous variables, for both the supply chain disrupted and undisrupted regimes. The solid black (solid red) line shows the point-wise posterior medians, and the shaded black area (dotted red lines) depicts the 68% equal-tailed point-wise posterior probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 10,000 independent draws from the posterior.

I. State-Dependence Results Using Local Projections

As a robustness check to our state-dependence results obtained using the TVAR model, we work with LPs to identify a contractionary monetary policy shock and analyze how it affects the macro aggregates for the U.S. economy depending on the level of global supply chain disruptions. LPs are a flexible approach that allows us to address the state-dependent effects of monetary policy without making strong parametric assumptions. Specifically, we use the LPs with interaction terms as in Ramey and Zubairy (2018), Ghassibe and Zanetti (2022), and Arias et al. (2023), and our identification scheme consists of sign restrictions implemented as described in Plagborg-Møller and Wolf (2021). Consider the following $n \times (K+1)$ projections:

$$\begin{aligned} y_{i,t+k} =& I_t \left[\boldsymbol{\beta}'_{\mathbb{D},i,k,0} \boldsymbol{y}_t + \sum_{l=1}^{L} \boldsymbol{\beta}'_{\mathbb{D},i,k,l} \boldsymbol{y}_{t-l} + \boldsymbol{C}'_{\mathbb{D},i,k} \boldsymbol{\omega}_t \right] \\ &+ (1 - I_t) \left[\boldsymbol{\beta}'_{\mathbb{U},i,k,0} \boldsymbol{y}_t + \sum_{l=1}^{L} \boldsymbol{\beta}'_{\mathbb{U},i,k,l} \boldsymbol{y}_{t-l} + \boldsymbol{C}'_{\mathbb{U},i,k} \boldsymbol{\omega}_t \right] + u_{i,k,t}, \end{aligned}$$

where $1 \leq i \leq n$, $0 \leq k \leq K$, $\boldsymbol{y}_t$ is an $n \times 1$ vector of the same endogenous variables as in Section 5.2 save for the ACR index (since it is the variable that we use to split the sample), $y_{i,t+k}$ is the value of the i-th variable in $\boldsymbol{y}_{t+k}$, $\boldsymbol{\omega}_t = [1, t]'$ is a 2×1 vector of a constant and a linear trend, and $u_{i,k,t}$ is the reduced-form error corresponding to the i-th variable. The vector of the reduced-form errors for $k = 1$, $\boldsymbol{u}_{1,t} = [u_{1,1,t} \quad \dots \quad u_{n,1,t}]'$, is assumed to have mean zero and variance-covariance matrix equal to $\mathbb{E}(\boldsymbol{u}_{1,t} \boldsymbol{u}'_{1,t}) = \boldsymbol{\Sigma}$.

Similar to the approach in the TVAR model, I_t serves as a dummy variable indicating whether the supply chain is disrupted. The regime of supply chain disruptions is defined based on whether the one-month lag of the ACR index exceeds its median level over the sample period. Figure I.1 illustrates the time series of the ACR index alongside its sample median. It is observed that, prior to mid-2017, the ACR index was consistently above its sample median, signifying the presence of the supply chain disrupted regime. However, from the second half of 2017 until the end of 2020, the ACR index remained below its sample median. This pattern reversed from early 2021 onward, with the ACR index rising above the median, indicating that the U.S. economy entered the disrupted regime. It is noteworthy that the transitions between the two regimes, as depicted in Figure I.1, closely align with those shown in Figure G.2, where the threshold $\overline{ACR}$ is determined endogenously in the TVAR estimation.

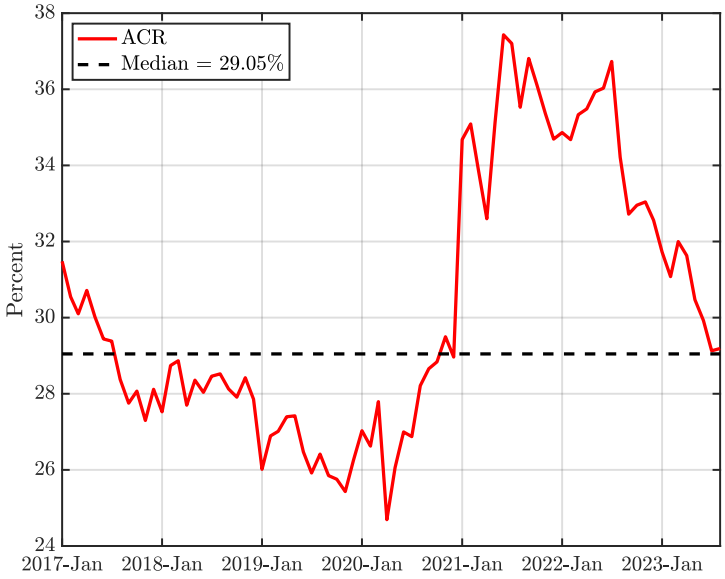

Figure I.1: ACR and Its Sample Median

Notes. The figure plots the ACR index as well as its sample median during the sample period from January 2017 to September 2023. The ACR index is computed using the AIS data of container ships and our IMA-DBSCAN algorithm, as detailed in Appendix B. The ACR index is measured in percentage terms and has been seasonally adjusted.

With the two regimes defined, the parameters $\boldsymbol{\beta}_{\mathbb{D},i,k,0}, \boldsymbol{\beta}_{\mathbb{D},i,k,l}$, and $\boldsymbol{C}_{\mathbb{D},i,k}$ correspond to the supply chain disrupted regime ($\mathbb{D}$), whereas the parameters $\boldsymbol{\beta}_{\mathbb{U},i,k,0}, \boldsymbol{\beta}_{\mathbb{U},i,k,l}$, and $\boldsymbol{C}_{\mathbb{U},i,k}$ correspond to the supply chain undisrupted regime ($\mathbb{U}$). As in our choice of the lag structure in the TVAR model, we include only one lag in the estimation of the LPs so as to reduce parameter uncertainty.[21]

In order to identify a contractionary monetary policy shock, we follow our theoretical prediction in Proposition 6 and come up with an identification scheme similar to that in Section 5.2. Yet, since the ACR index is not included in the estimation, we drop the zero restriction in Restriction 4 and re-write it as the following:

Restriction 4′. ***A contractionary monetary policy shock*** *leads to a negative response of real GDP, PCE goods price, product market tightness, and import price, as well as to a positive response of spare capacity and the federal funds rate at* $k = 1, 2, 3$. *In addition, the on-impact response of spare capacity in p.p. is bounded to be smaller than one hundred times that of the federal funds rate in p.p.*

Restriction 4′ is similar to Restriction 4, except that we impose restrictions on the subsequent

[21] We have also tried two lags in the estimation and derived similar results in Figure I.5.

horizons to sharpen our identification, and an elasticity bound is imposed to discipline the identified set of IRFs corresponding to spare capacity. The latter variation is critical to ensure that our estimation is plausible, since in the absence of such a bound, the identified set of IRFs would include an increase in the spare capacity rate of 100 p.p. as being equally likely as an increase in the spare capacity rate of 1 p.p., following an unexpected increase in the federal funds rate of 0.05 p.p. Hence, we use the elasticity bound to rule out dubious IRFs following Kilian and Murphy (2012) and Arias et al. (2019, 2023).

With Restriction 4′, we determine the identified set of IRFs for each regime by numerically solving the quadratic program outlined in the supplement to Plagborg-Møller and Wolf (2021), using Algorithm 2 from Giacomini and Kitagawa (2021). Without loss of generality, we normalize the first shock to be the shock of interest. Let $\boldsymbol{S}_1$ denote a $12 \times n$ matrix that selects the IRFs that we restrict to be negative, and let $\boldsymbol{S}_2$ denote a $6 \times n$ matrix that selects the IRFs that we restrict to be positive (there are a total of 18 sign restrictions in Restriction 4′). Then, for each regime, we draw $D = 100,000$ orthogonal matrices $\boldsymbol{Q}_{r,d}$ (i.e., $\boldsymbol{Q}'_{r,d}\boldsymbol{Q}_{r,d} = \boldsymbol{Q}_{r,d}\boldsymbol{Q}'_{r,d} = \mathbf{1}_{n\times n}$) that satisfy the following:

$$\begin{aligned} \boldsymbol{S}_1\hat{\boldsymbol{B}}_{r,0:2}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1 &\leq 0, \\ \boldsymbol{S}_2\hat{\boldsymbol{B}}_{r,0:2}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1 &\geq 0, \\ \frac{\boldsymbol{e}_4'\hat{\boldsymbol{B}}_{r,0}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1}{\boldsymbol{e}_1'\hat{\boldsymbol{B}}_{r,0}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1} - 100 &\leq 0, \end{aligned} \tag{I.1}$$

where $r \in \{\mathbb{D}, \mathbb{U}\}$, $1 \leq d \leq D$, $\hat{\boldsymbol{B}}_{r,0:2} = [\hat{\boldsymbol{B}}'_{r,0} \;\; \hat{\boldsymbol{B}}'_{r,1} \;\; \hat{\boldsymbol{B}}'_{r,2}]'$, $\hat{\boldsymbol{B}}_{r,k} = [\hat{\boldsymbol{\beta}}_{r,1,k,0} \;\; \dots \;\; \hat{\boldsymbol{\beta}}_{r,n,k,0}]'$, $\hat{\boldsymbol{\beta}}_{r,i,k,0}$ is the OLS estimate of $\boldsymbol{\beta}_{r,i,k,0}$, $\hat{\boldsymbol{\Omega}} = chol(\hat{\boldsymbol{\Sigma}})'$, $chol$ is the upper triangular Cholesky decomposition of $\hat{\boldsymbol{\Sigma}}$, and $\hat{\boldsymbol{\Sigma}}$ is the OLS estimate of $\boldsymbol{\Sigma}$.[22] Given that the entry (i, j) in $\hat{\boldsymbol{B}}_{r,k}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}$ gives the response of the i-th endogenous variable to the j-th shock at horizon k, the first two inequality conditions in Equation (I.1) summarize all the sign restrictions imposed on the IRFs, while the last inequality condition contains the elasticity bound, as $(\boldsymbol{e}_4'\hat{\boldsymbol{B}}_{r,0}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1)/(\boldsymbol{e}_1'\hat{\boldsymbol{B}}_{r,0}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1)$ denotes the ratio between the on-impact responses of spare capacity and the federal funds rate, where e_i is the i-th column of the n-dimensional identity matrix.

Given $\hat{\boldsymbol{B}}_{r,k}$ and $\hat{\boldsymbol{\Omega}}$, let $\{\boldsymbol{Q}_{r,d}\}_{d=1,\dots,D}$ be the draws of orthogonal matrices that satisfy the restrictions in Equation (I.1). The identified set of IRFs of the i-th endogenous variable at

[22] Vector inequalities are to be understood element-wise.

horizon k is thus given by:

$$\left[\min_{d}\left\{0.05\frac{\boldsymbol{e}_i'\hat{\boldsymbol{B}}_{r,k}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1}{\boldsymbol{e}_1'\hat{\boldsymbol{B}}_{r,0}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1}\right\}_{d=1,\ldots,D},\ \max_{d}\left\{0.05\frac{\boldsymbol{e}_i'\hat{\boldsymbol{B}}_{r,k}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1}{\boldsymbol{e}_1'\hat{\boldsymbol{B}}_{r,0}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1}\right\}_{d=1,\ldots,D}\right], \tag{I.2}$$

where the factor $0.05/(\boldsymbol{e}_1'\hat{\boldsymbol{B}}_{r,0}\hat{\boldsymbol{\Omega}}\boldsymbol{Q}_{r,d}\boldsymbol{e}_1)$ is a normalization imposed so that in both regimes, the contractionary monetary policy shock raises the federal funds rate by 0.05 p.p. on impact.

Figure I.2 presents the point-wise medians and the 68% equal-tailed point-wise probability bands for the identified set of IRFs in each regime following a contractionary monetary policy shock. In line with the TVAR model, the IRFs are shown from horizon $k = 0$ to horizon $k = 12$, which equates to four quarters. As depicted in Figure I.2, the state-dependent effects of a contractionary monetary policy shock are still evident. More specifically, the responses of real GDP and spare capacity are more subdued. At the same time, those of the PCE goods and import prices are more pronounced during periods of global supply chain disruption. However, the differences in the responses of the federal funds rate and product market tightness between the two regimes are not distinctly discernible.

We have also explored alternative thresholds at the 40th and 60th percentiles of the ACR index to differentiate between the supply chain disrupted and undisrupted regimes. The findings, as illustrated in Figures I.3 and I.4, demonstrate that our main results remain robust under these varying thresholds.

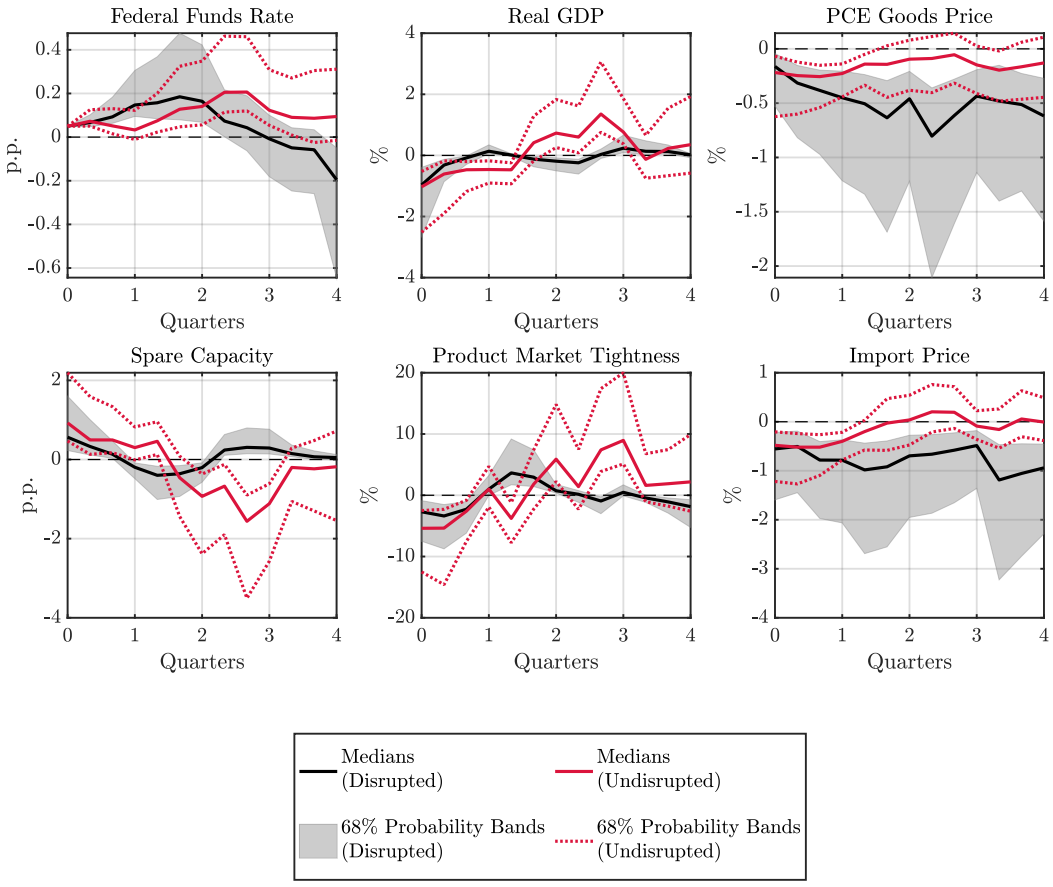

Figure I.2: State-Dependent Effects of a Contractionary Monetary Policy Shock: Using the LPs With Interaction Terms and a Threshold at the Median of the ACR Index

Notes. The figure shows the IRFs to a contractionary monetary policy shock identified using the LPs with interaction terms as in Ramey and Zubairy (2018), Ghassibe and Zanetti (2022), and Arias et al. (2023), along with Restriction 4′ imposed on the IRFs of the endogenous variables, for both the supply chain disrupted and undisrupted regimes. A threshold at the sample median of the ACR index is applied to distinguish between the two regimes. The solid black (solid red) line shows the point-wise medians and the shaded black area (dotted red lines) shows the 68% equal-tailed point-wise probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 100,000 draws of orthogonal matrices.

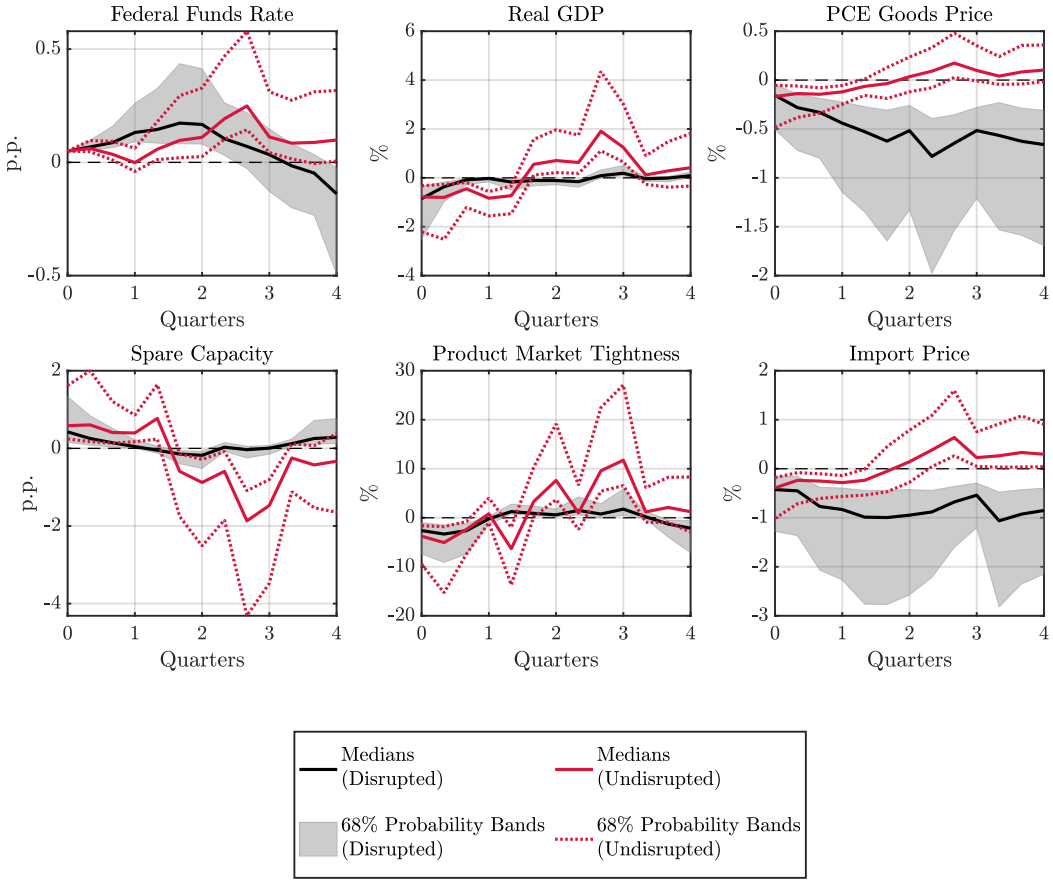

Figure I.3: State-Dependent Effects of a Contractionary Monetary Policy Shock: Using the LPs With Interaction Terms and a Threshold at the 40th Percentile of the ACR Index

Notes. The figure shows the IRFs to a contractionary monetary policy shock identified using the LPs with interaction terms as in Ramey and Zubairy (2018), Ghassibe and Zanetti (2022), and Arias et al. (2023), along with Restriction 4′ imposed on the IRFs of the endogenous variables, for both the supply chain disrupted and undisrupted regimes. A threshold at the 40th percentile of the ACR index is applied to distinguish between the two regimes. The solid black (solid red) line shows the point-wise medians and the shaded black area (dotted red lines) shows the 68% equal-tailed point-wise probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 100,000 draws of orthogonal matrices.

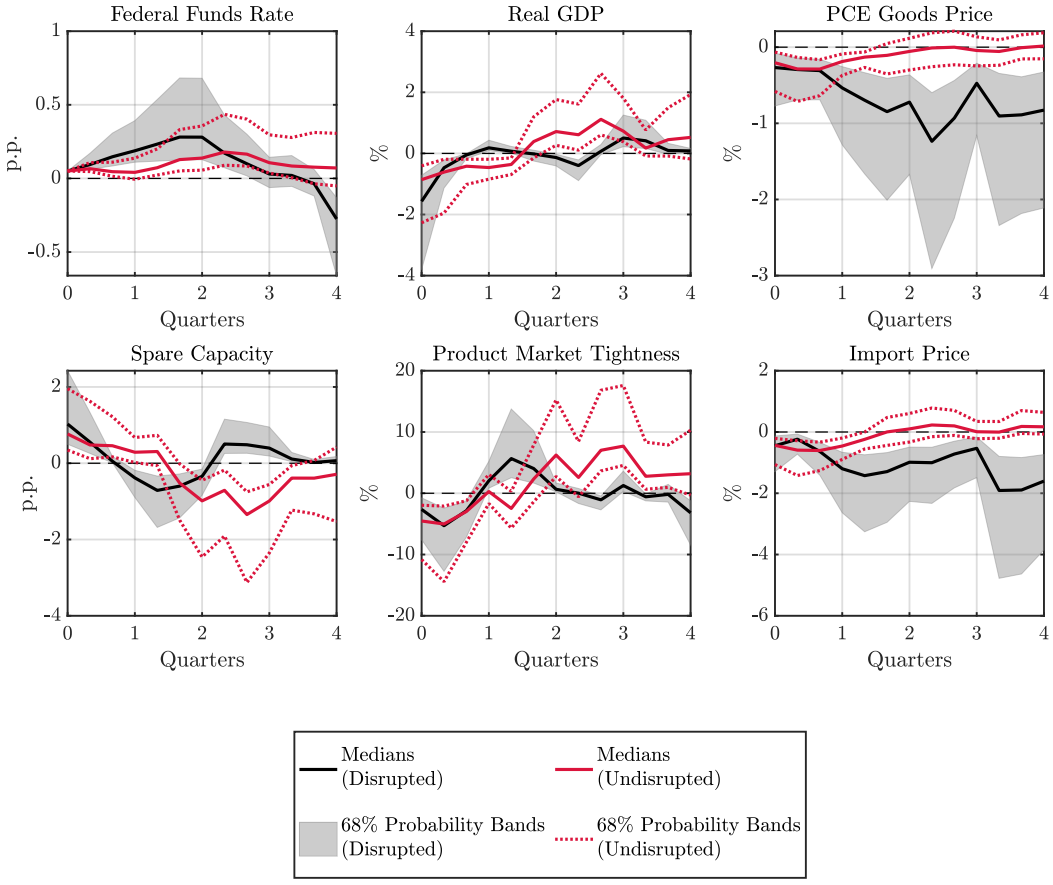

Figure I.4: State-Dependent Effects of a Contractionary Monetary Policy Shock: Using the LPs With Interaction Terms and a Threshold at the 60th Percentile of the ACR Index

Notes. The figure shows the IRFs to a contractionary monetary policy shock identified using the LPs with interaction terms as in Ramey and Zubairy (2018), Ghassibe and Zanetti (2022), and Arias et al. (2023), along with Restriction 4′ imposed on the IRFs of the endogenous variables, for both the supply chain disrupted and undisrupted regimes. A threshold at the 60th percentile of the ACR index is applied to distinguish between the two regimes. The solid black (solid red) line shows the point-wise medians and the shaded black area (dotted red lines) shows the 68% equal-tailed point-wise probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 100,000 draws of orthogonal matrices.

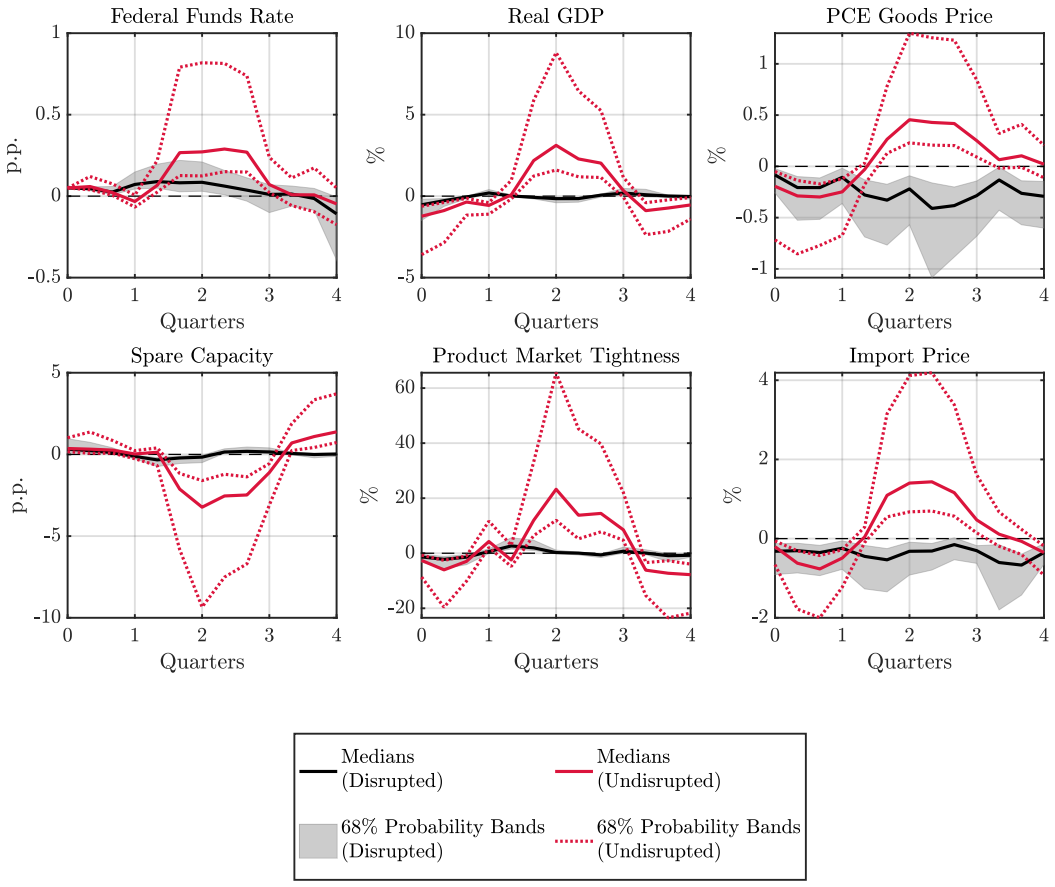

Figure I.5: State-Dependent Effects of a Contractionary Monetary Policy Shock: Using the LPs With Interaction Terms and Two Lags in the Estimation

Notes. The figure shows the IRFs to a contractionary monetary policy shock identified using the LPs with interaction terms as in Ramey and Zubairy (2018), Ghassibe and Zanetti (2022), and Arias et al. (2023), along with Restriction 4′ imposed on the IRFs of the endogenous variables, for both the supply chain disrupted and undisrupted regimes. A threshold at the sample median of the ACR index is applied to distinguish between the two regimes, and two lags are included in the estimation. The solid black (solid red) line shows the point-wise medians and the shaded black area (dotted red lines) shows the 68% equal-tailed point-wise probability bands for the supply chain disrupted (undisrupted) regime. The figure is based on 100,000 draws of orthogonal matrices.

References for Appendices

Adland, R. and Jia, H. (2016). Dynamic Speed Choice in Bulk Shipping. *Maritime Economics & Logistics*, 20(3):253–266.

Allen, T. (2014). Information Frictions in Trade. *Econometrica*, 82:2041–2083.

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

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

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

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

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics*, 131:1593–1636.

Barnichon, R. and Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *The Review of Economics and Statistics*, 101(3):522–530.

Bańbura, M., Giannone, D., and Reichlin, L. (2010). Large Bayesian Vector Auto Regressions. *Journal of Applied Econometrics*, 25:71–92.

Benguria, F. (2021). The Matching and Sorting of Exporting and Importing Firms: Theory and Evidence. *Journal of International Economics*, 131:103430.

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

Bils, M., Chang, Y., and Kim, S.-B. (2011). Worker Heterogeneity and Endogenous Separations in a Matching Model of Unemployment Fluctuations. *American Economic Journal: Macroeconomics*, 3:128–154.

Birant, D. and Kut, A. (2007). ST-DBSCAN: An Algorithm for Clustering Spatial–Temporal Data. *Data & Knowledge Engineering*, 60:208–221.

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

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

Burriel, P., Kataryniuk, I., Pérez, C. M., and Viani, F. (2023). A New Supply Bottlenecks Index Based on Newspaper Data. *Banco de España Working Paper 2304*.

Chaney, T. (2014). The Network Structure of International Trade. *American Economic Review*, 104(11):3600–3634.

Chow, G. C. and Lin, A.-l. (1971). Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series. *Review of Economics and Statistics*, 53:372–375.

di Giovanni, J., Şebnem Kalemli-Özcan, Silva, A., and Yildirim, M. A. (2022). Global Supply Chain Pressures, International Trade, and Inflation. *NBER Working Paper 30240*.

Du, Y., Chen, Q., Lam, J. S. L., Xu, Y., and Cao, J. X. (2015). Modeling the Impacts of Tides and the Virtual Arrival Policy in Berth Allocation. *Transportation Science*, 49(4):939–956.

Eaton, J. and Kortum, S. (2002). Technology, Geography, and Trade. *Econometrica*, 70:1741–1779.

Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases With Noise. *KDD-96 Proceedings*.

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

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

Finck, D. and Tillmann, P. (2022). The Macroeconomic Effects of Global Supply Chain Disruptions. *BOFIT Discussion Paper 14/2022*.

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

Fujita, S. and Ramey, G. (2012). Exogenous Versus Endogenous Separation. *American Economic Journal: Macroeconomics*, 4:68–93.

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

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

Jordà, . (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

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

Kilian, L. and Murphy, D. P. (2012). Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models. *Journal of the European Economic Association*, 10:1166–1188.

Kilian, L., Nomikos, N., and Zhou, X. (2023). Container Trade and the U.S. Recovery. *International Journal of Central Banking*, 19(1):417–450.

Krolikowski, P. M. and McCallum, A. H. (2021). Goods-Market Frictions and International Trade. *Journal of International Economics*, 129:103411.

Lenoir, C., Martin, J., and Mejean, I. (2022). Search Frictions in International Goods Markets. *Journal of the European Economic Association*, 21(1):326–366.

Li, C., Qi, X., and Song, D. (2016). Real-Time Schedule Recovery in Liner Shipping Service With Regular Uncertainties and Disruption Events. *Transportation Research Part B: Methodological*, 93:762–788.

Litterman, R. B. (1986). Forecasting With Bayesian Vector Autoregressions: Five Years of Experience. *Journal of Business & Economic Statistics*, 4:25–38.

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

Melitz, J. and Toubal, F. (2014). Native Language, Spoken Language, Translation and Trade. *Journal of International Economics*, 93:351–363.

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, 71:1695–1725.

Menzio, G. and Shi, S. (2011). Efficient Search on the Job and the Business Cycle. *Journal of Political Economy*, 119:468–510.

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

Mountford, A. and Uhlig, H. (2009). What Are the Effects of Fiscal Policy Shocks? *Journal of Applied Econometrics*, 24:960–992.

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

Naudé, W. and Matthee, M. (2011). The Impact of Transport Costs on New Venture Internationalisation. *Journal of International Entrepreneurship*, 9:62–89.

Newey, W. K. and West, K. D. (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–708.

Notteboom, T. E. (2006). The Time Factor in Liner Shipping Services. *Maritime Economics & Logistics*, 8:19–39.

Pizzinelli, C., Theodoridis, K., and Zanetti, F. (2020). State Dependence in Labor Market Fluctuations. *International Economic Review*, 61:1027–1072.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica*, 89:955–980.

Ramey, V. A. and Zubairy, S. (2018). Government Spending Multipliers in Good Times and in Bad: Evidence From US Historical Data. *Journal of Political Economy*, 126:850–901.

Rodrigue, J.-P. (2020). *The Geography of Transport Systems*. Routledge, 5 edition.

Smirnyagin, V. and Tsyvinski, A. (2022). Macroeconomic and Asset Pricing Effects of Supply Chain Disasters. *NBER Working Paper 30503*.

Stamer, V. (2021). Thinking Outside the Container: A Sparse Partial Least Squares Approach to Forecasting Trade Flows. *Kiel Working Papers 2179*.

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

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

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

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

Wong, W. F. (2022). The Round Trip Effect: Endogenous Transport Costs and International Trade. *American Economic Journal: Applied Economics*, 14:127–166.

Wu, J. C. and Xia, F. D. (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking*, 48(2-3):253–291.