Supplemental Appendix

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

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

Contents

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

B Testing Short-Run Rigidity in Containerized Shipping **A-5**

C A Density-Based Spatial Clustering Algorithm **A-7**

C.1 Methodology A-7

C.2 Illustrative Cases A-15

C.3 Weekly Indices of Port Congestion A-22

D Discussion of Model Assumptions **A-24**

E Proofs and Discussions Omitted from the Main Text **A-26**

E.1 Proof of Proposition 1 . . . . . A-26

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

E.3 Slope and Curvature of Match Separation and Creation Schedules A-31

E.4 Proof of Comparative Statics in Table 1 . . . . . A-33

E.5 Convergence Dynamics A-36

F External Data Sources for Baseline Estimation **A-42**

G Robustness of Baseline Results **A-44**

G.1 Prior Robustness A-44

G.2 Price Levels vs. Inflation Rates in the SVAR A-46

G.3 Invertibility Under an Expanded Information Set A-51

† Bai: Tsinghua University, China. xiwenbai@mail.tsinghua.edu.cn. ‡ Fernández-Villaverde: University of Pennsylvania, U.S. jesusfv@econ.upenn.edu. § Li: Peking University HSBC Business School & UIBE, China. yiliang_li@outlook.com. ¶ Zanetti: University of Oxford, U.K. francesco.zanetti@economics.ox.ac.uk.

G.4 Nonlinear Transmission by Shock Size A-54
G.5 Different Lag Structures A-57
G.6 Dropping the Linear Trend A-61
G.7 Alternative Proxies for Activity and Prices A-63

H Identification Gains From Supply Chain Proxies **A-73**

H.1 No Proxy A-73
H.2 ACR Without Zero Restrictions A-76
H.3 ACT A-78
H.4 Trans-Pacific ACR A-80
H.5 HARPEX A-83
H.6 GSCPI A-86
H.7 Supply Disruptions Index (SDI) A-89
H.8 Goods Price Responses Across Proxies A-92

Appendix References **A-93**

Data References **A-96**

A. Background on the Containerized Shipping Industry

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

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

Oil price, vessel speed, and port congestion. Fuel costs account for 50–60% of a liner shipping company's operating expenses (Notteboom, 2006), and vessel fuel consumption increases roughly with the cube of sailing speed (Li, Qi and Song, 2016). As a result, container ship operators adjust sailing speeds in response to fluctuations in bunker oil prices. We estimate a strong negative relationship between oil prices and average vessel speed: a 1% increase in the Brent crude price reduces sailing speed by about 0.022%, a highly significant effect ($p = 0.004$). In contrast, when we regress congestion rates at the top 50 container ports worldwide on Brent crude prices (controlling for year and month fixed effects), we find no statistically significant relationship: the estimated coefficient is positive but small relative to its standard error, with a p-value of 0.313 and an R^2 of only 0.021.[1]

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

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

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

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

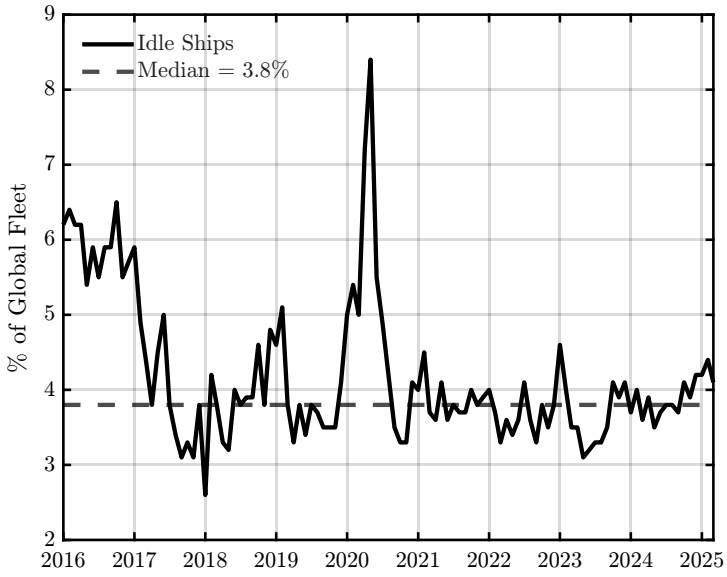


Figure A.1: Share of Idle Container Ships in the Global Fleet

B. Testing Short-Run Rigidity in Containerized Shipping

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

Our assessment involves estimating

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

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

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

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

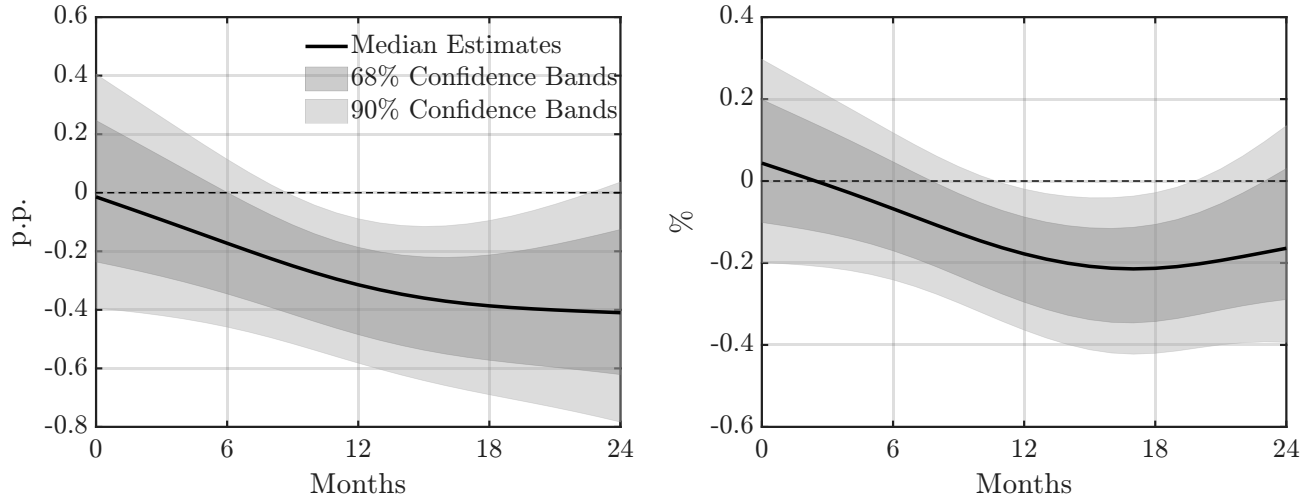
where $B_m(k)$ are B-spline basis functions and b_m are parameters estimated using a penalized procedure that shrinks Equation (B.2) toward a quadratic polynomial in k.

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

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

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

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



(a) ACR **(b)** ACT

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

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

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

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

C. A Density-Based Spatial Clustering Algorithm

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

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

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

C.1. Methodology

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

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

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

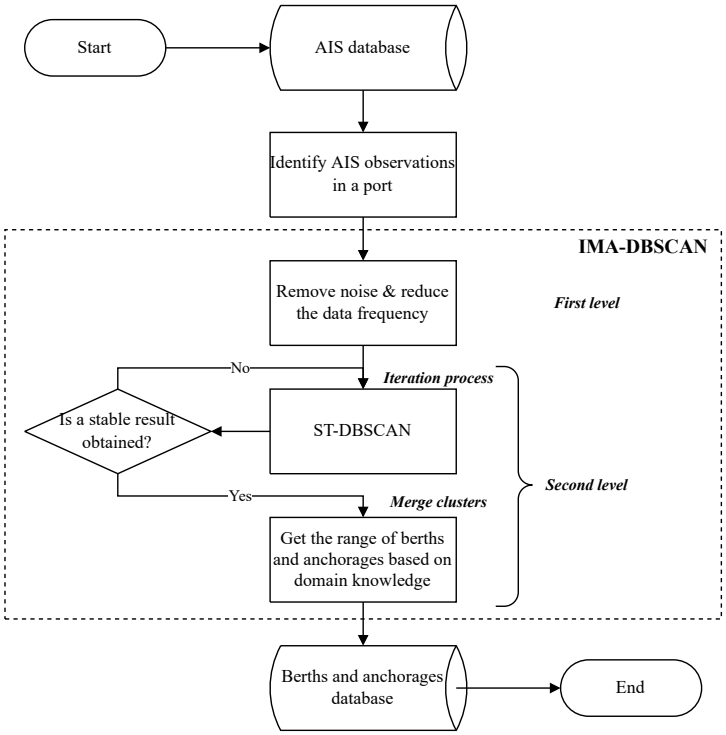


Figure C.1: Methodology Framework of IMA-DBSCAN

C.1.1. First Level: Data Preprocessing

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

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

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

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

C.1.2. Second Level: Multiple Attributes and Iteration

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

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

The geographical distance D is computed using the Haversine formula:

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

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

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

spatial distance Δh between two headings is defined as:

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

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

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

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

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

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

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

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

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

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

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

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

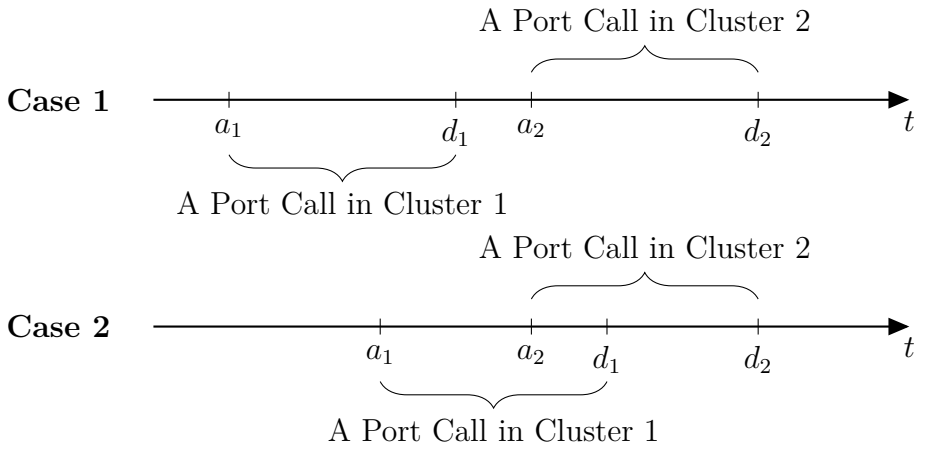


Figure C.2: Merging Clusters

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

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

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

Table C.1: Parameter Values for IMA-DBSCAN

	Parameter					
First Level	Δt	Eps	$MinPts$			
Value	12 hours	50 m	10			
Second Level	α	β	$\Delta Dist$	Eps'	$MinPts'$	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 heading 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 heading 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 (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 (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* (C.1)
    $Dist2 \leftarrow \Delta h(o, o')$ ▷ *See Equation* (C.2)
    if $Dist1 \leq Eps1$ **and** $Dist2 \leq Eps2$ **then**
      Append o' to $Neighbors$
    end if
  end for
  return $Neighbors$
end function

C.2. Illustrative Cases

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

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

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

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

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

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

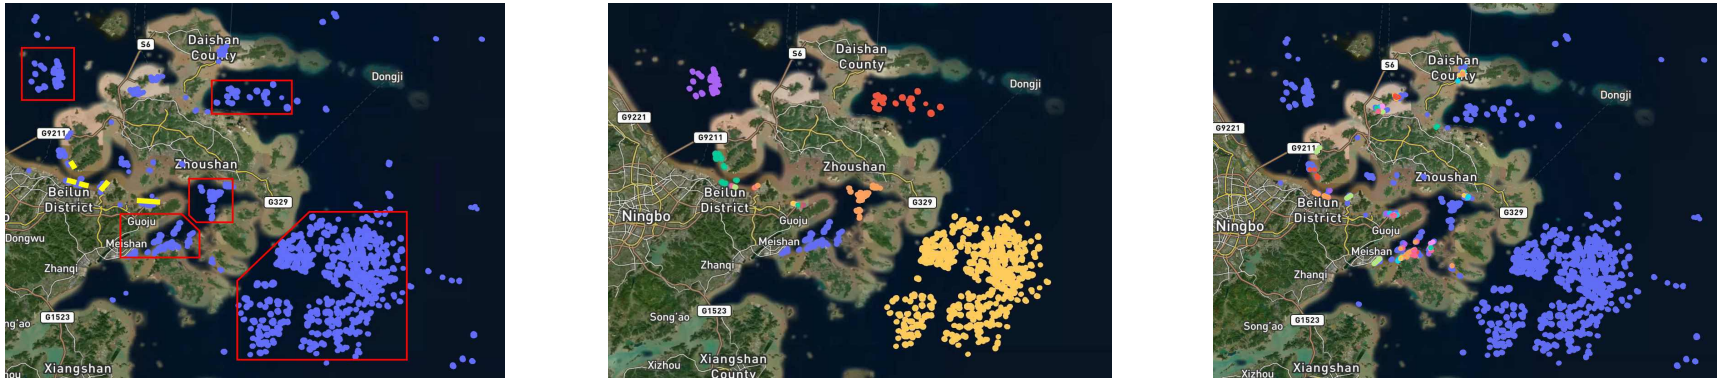
To benchmark IMA-DBSCAN against existing methods, we compare it with ST-DBSCAN.[10] ST-DBSCAN is widely recognized in the literature as a leading spatial clustering algorithm capable of handling spatio-temporal data. Figure C.3c reports its results, which are less precise than those of IMA-DBSCAN. Although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

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

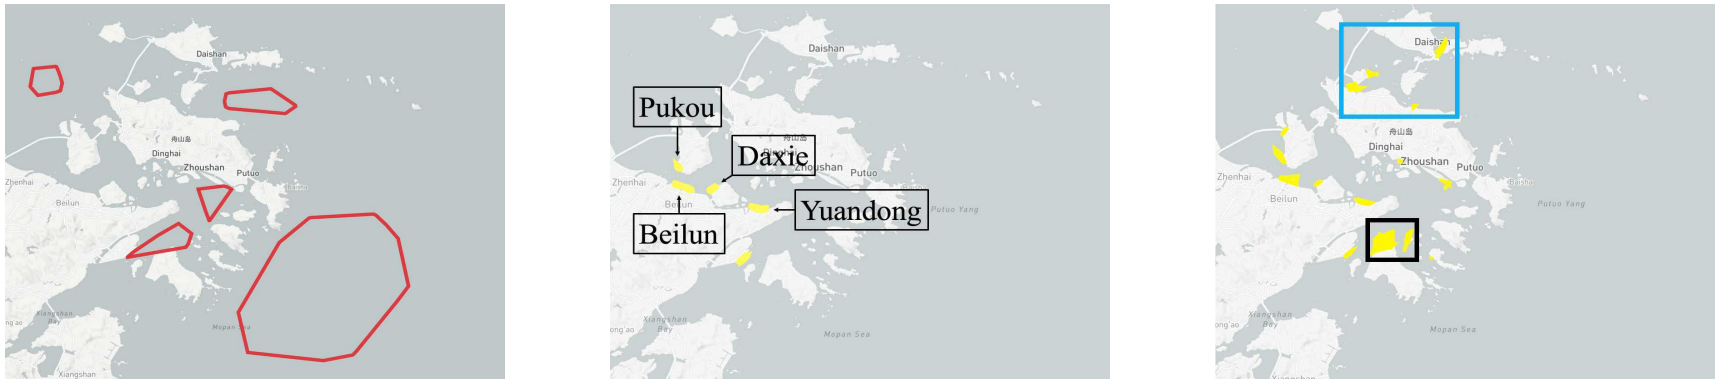
Finally, we apply both algorithms to the Ports of Los Angeles and Long Beach (U.S.), Rotterdam (Netherlands), and Singapore. As shown in Figures C.5 through C.7, IMA-DBSCAN consistently outperforms ST-DBSCAN, delivering more accurate identification across all major container ports examined.

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

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



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



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

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

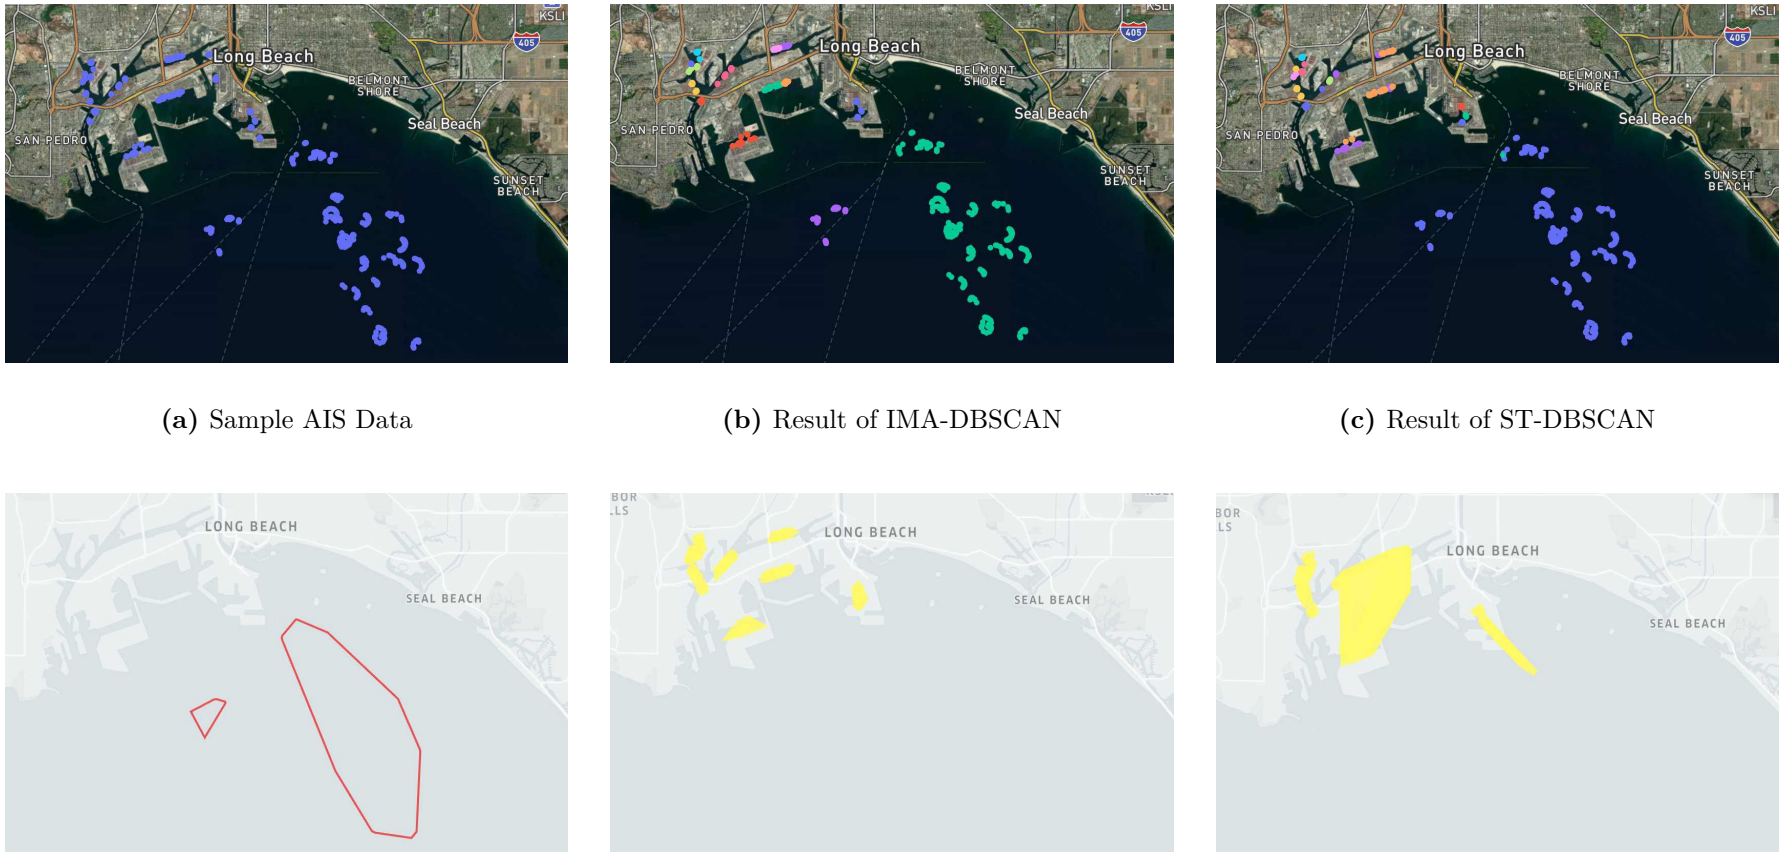
Notes. In Panel (a), the sample data consist of the first 50,000 AIS observations recorded since January 1, 2020, within the Port of Ningbo-Zhoushan in China. These are shown as blue dots, covering coordinates from 121.60°E to 123.00°E and from 29.50°N to 30.35°N. As benchmarks, satellite maps and nautical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



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

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

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



(a) Sample AIS Data

(b) Result of IMA-DBSCAN

(c) Result of ST-DBSCAN

(d) Anchorages (IMA-DBSCAN)

(e) Berths (IMA-DBSCAN)

(f) Berths (ST-DBSCAN)

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

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



(a) Sample AIS Data

(b) Result of IMA-DBSCAN

(c) Result of ST-DBSCAN

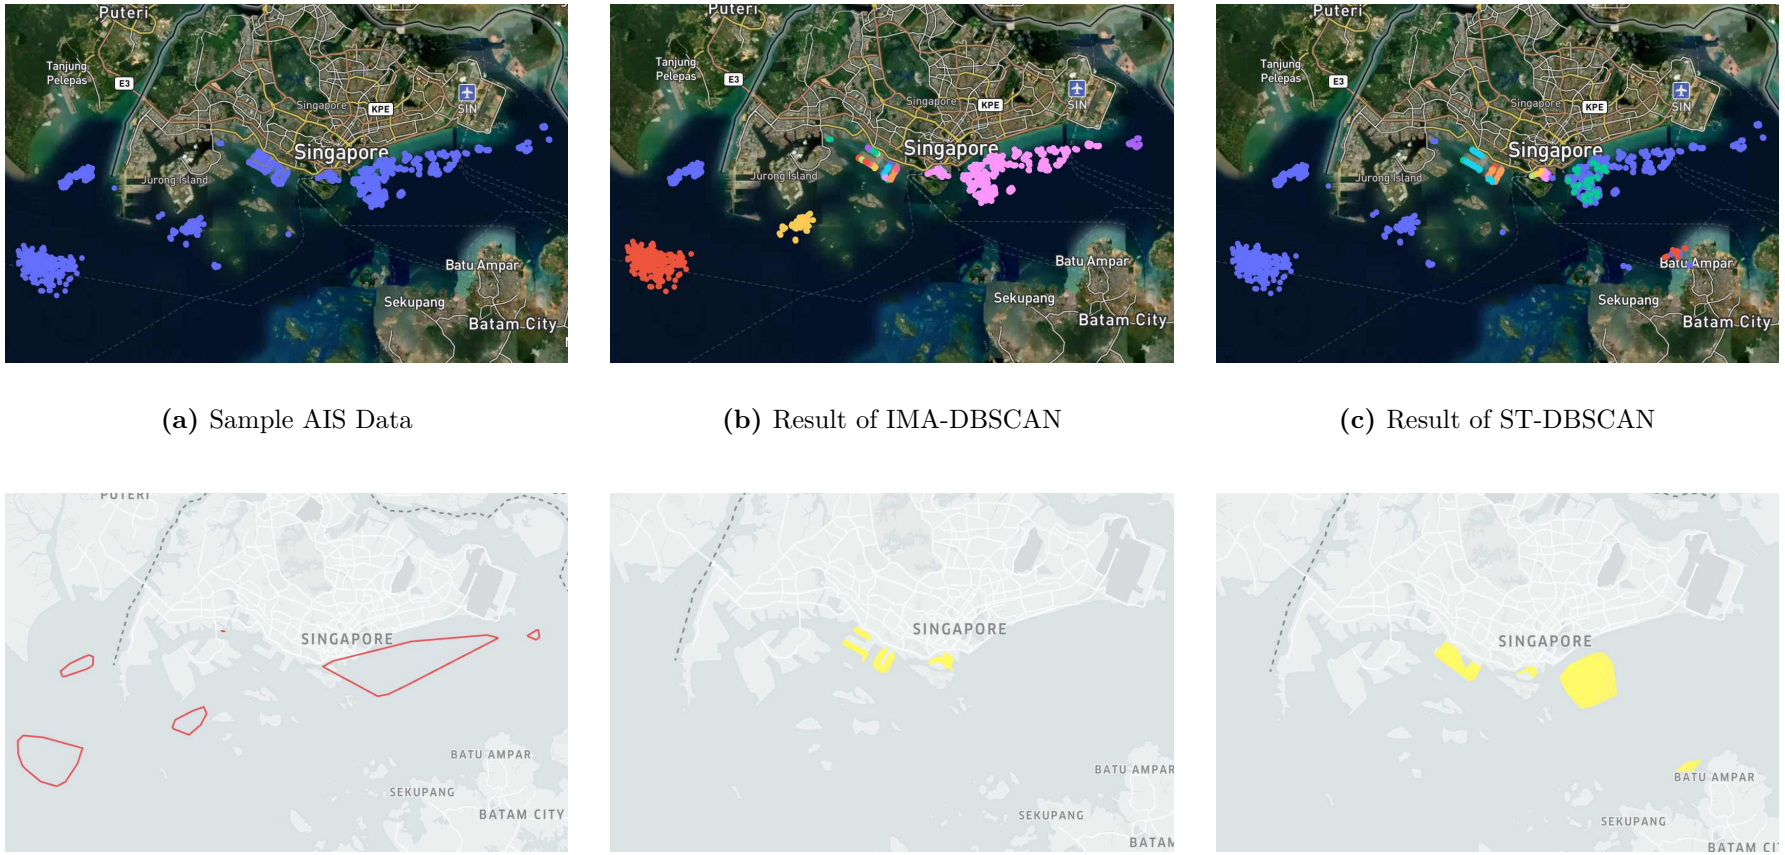
(d) Anchorages (IMA-DBSCAN)

(e) Berths (IMA-DBSCAN)

(f) Berths (ST-DBSCAN)

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

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



(a) Sample AIS Data

(b) Result of IMA-DBSCAN

(c) Result of ST-DBSCAN

(d) Anchorages (IMA-DBSCAN)

(e) Berths (IMA-DBSCAN)

(f) Berths (ST-DBSCAN)

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

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

C.3. Weekly Indices of Port Congestion

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

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

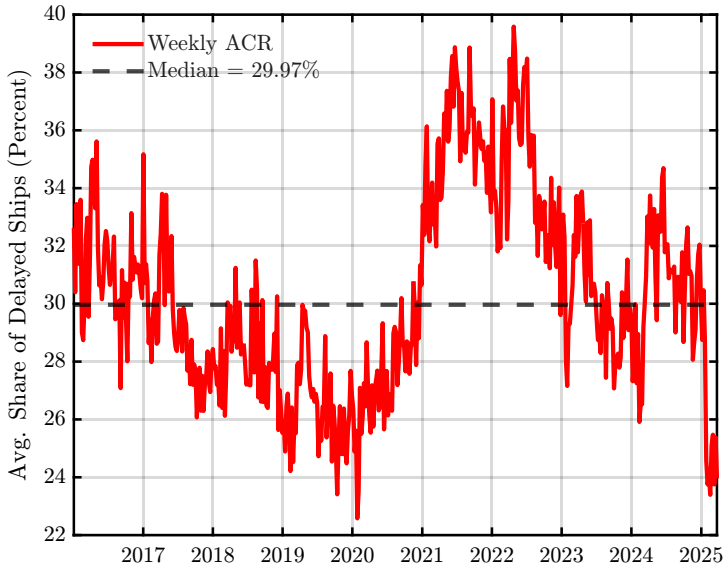


Figure C.8: Weekly ACR Index

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

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

worldwide:

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

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

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

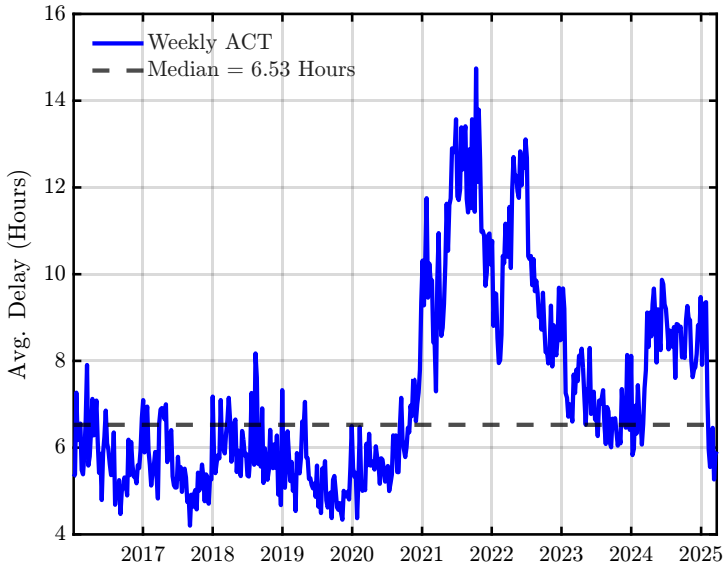


Figure C.9: Weekly ACT Index

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

D. Discussion of Model Assumptions

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

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

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

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

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

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

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

E. Proofs and Discussions Omitted from the Main Text

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

E.1. Proof of Proposition 1

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

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

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

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

where the second line follows by integration by parts.

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

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

and the rearranged match creation condition:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

By the chain rule and the signs above,

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

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

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

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

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

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

E.2. Proof of Proposition 2

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

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

and the rearranged match creation condition is:

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

where

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

and hence

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

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

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

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

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

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

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

we obtain:

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

which is locally convex.

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

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

if $\xi = 1$,

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

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

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

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

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

Define:

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

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

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

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

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

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

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

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

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

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

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

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

Differentiating once more gives:

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

where the partial derivatives are:

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

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

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

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

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

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

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

which is well-defined on the interior domain:

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

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

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

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

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

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

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

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

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

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

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

E.4. Proof of Comparative Statics in Table 1

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

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

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

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

the Jacobian with respect to x is given by:

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

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

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

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

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

by the Implicit Function Theorem,

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

and Cramer's rule yields:

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

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

Given that:

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

we obtain:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Combine the pieces:

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

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

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

Therefore, the closed form of r_γ is given by:

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

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

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

E.5. Convergence Dynamics

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

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

Three equations govern these variables:

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

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

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

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

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

Table E.1: Baseline Calibration

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

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

E.5.1. An Adverse Shock to Aggregate Demand

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

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

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

Iteration	# Matched Producers	Consumption (or Output)	Price	Reservation Transportation Cost	Product Market Tightness	Wholesale Price	Spare Capacity
	x_M	c	p	$\bar{z}$	θ	r	$l - c$
Initial Steady State	0.8281	0.8281	12.0753	18.1291	12.0247	18.1084	0.1719
1	0.8206	0.8206	11.5773	17.1049	11.1775	17.1355	0.1794
2	0.8219	0.8219	11.5584	17.2839	11.3250	17.2524	0.1781
3	0.8217	0.8217	11.5617	17.2518	11.2986	17.2314	0.1783
New Steady State	0.8217	0.8217	11.5612	17.2567	11.3026	17.2346	0.1783

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

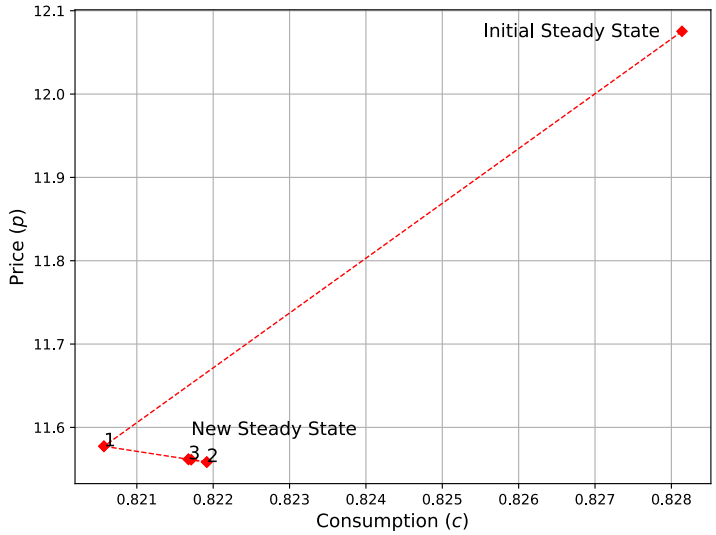


Figure E.1: Convergence Dynamics of Consumption and Price: An Adverse Shock to Aggregate Demand ($\mu = 10 \to 9.5$)

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

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

Iteration	# Matched Producers	Consumption (or Output)	Price	Reservation Transportation Cost	Product Market Tightness	Wholesale Price	Spare Capacity
	x_M	c	p	$\bar{z}$	θ	r	$l - c$
Initial Steady State	0.8281	0.8281	12.0753	18.1291	12.0247	18.1084	0.1719
1	0.8128	0.8128	11.1038	16.1385	10.3850	16.2174	0.1872
2	0.8156	0.8156	11.0655	16.4846	10.6680	16.4420	0.1844
3	0.8151	0.8151	11.0726	16.4202	10.6153	16.4002	0.1849
New Steady State	0.8152	0.8152	11.0715	16.4301	10.6234	16.4067	0.1848

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

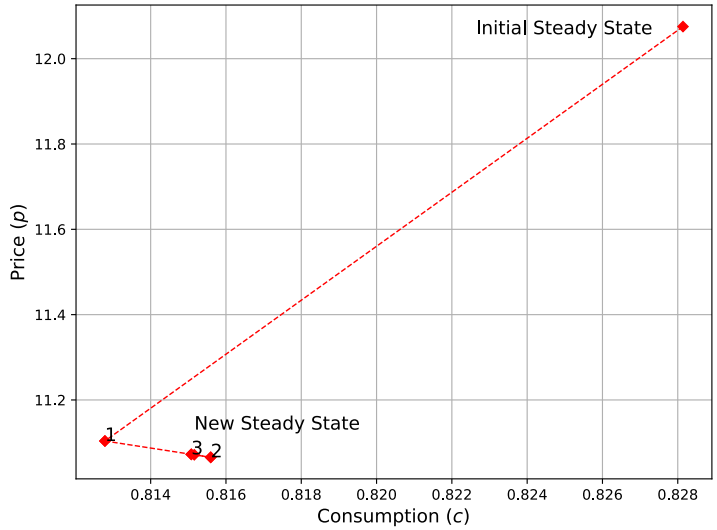


Figure E.2: Convergence Dynamics of Consumption and Price: An Adverse Shock to Aggregate Demand ($\chi = 1 \to 0.95$)

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

E.5.2. An Adverse Shock to Productive Capacity

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

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

Iteration	# Matched Producers	Consumption (or Output)	Price	Reservation Transportation Cost	Product Market Tightness	Wholesale Price	Spare Capacity
	x_M	c	p	$\bar{z}$	θ	r	$l - c$
Initial Steady State	0.8281	0.8281	12.0753	18.1291	12.0247	18.1084	0.1719
1	0.8355	0.7937	12.5989	19.2144	12.9300	19.1392	0.1645
2	0.8343	0.7925	12.6176	19.0229	12.7697	19.0126	0.1657
3	0.8345	0.7927	12.6144	19.0547	12.7964	19.0337	0.1655
New Steady State	0.8344	0.7927	12.6149	19.0502	12.7925	19.0307	0.1656

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

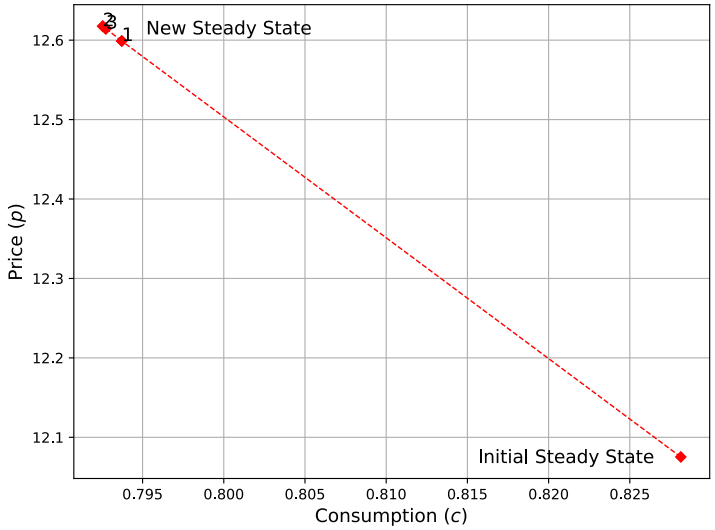


Figure E.3: Convergence Dynamics of Consumption and Price: An Adverse Shock to Productive Capacity ($l = 1 \to 0.95$)

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

E.5.3. An Adverse Shock to the Supply Chain

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

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

Iteration	# Matched Producers x_M	Consumption (or Output) c	Price p	Reservation Transportation Cost $\bar{z}$	Product Market Tightness θ	Wholesale Price r	Spare Capacity $l - c$
Initial Steady State	0.8281	0.8281	12.0753	18.1291	12.0247	18.1084	0.1719
1	0.8208	0.8208	12.1834	18.0045	11.7737	18.0374	0.1792
2	0.8221	0.8221	12.1640	18.1875	11.9245	18.1569	0.1779
3	0.8219	0.8219	12.1675	18.1547	11.8974	18.1354	0.1781
New Steady State	0.8219	0.8219	12.1669	18.1596	11.9015	18.1387	0.1781

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

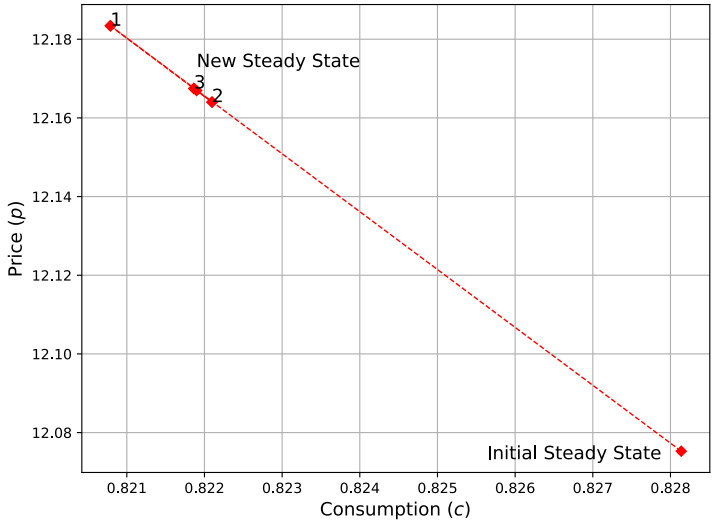


Figure E.4: Convergence Dynamics of Consumption and Price: An Adverse Shock to the Supply Chain ($\gamma = 1 \to 1.05$)

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

F. External Data Sources for Baseline Estimation

Table F.1 lists the external data sources used in the baseline SVAR estimation in Section III, together with any adjustments applied. Below, we provide further details on the construction of spare capacity and product market tightness. The data used for robustness checks are described in subsequent appendices. To conserve space, full reference entries for all external data sources are placed at the end of the Supplemental Appendix.

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

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

Table F.1: External Data Sources for Baseline Estimation

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

G. Robustness of Baseline Results

G.1. Prior Robustness

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

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

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

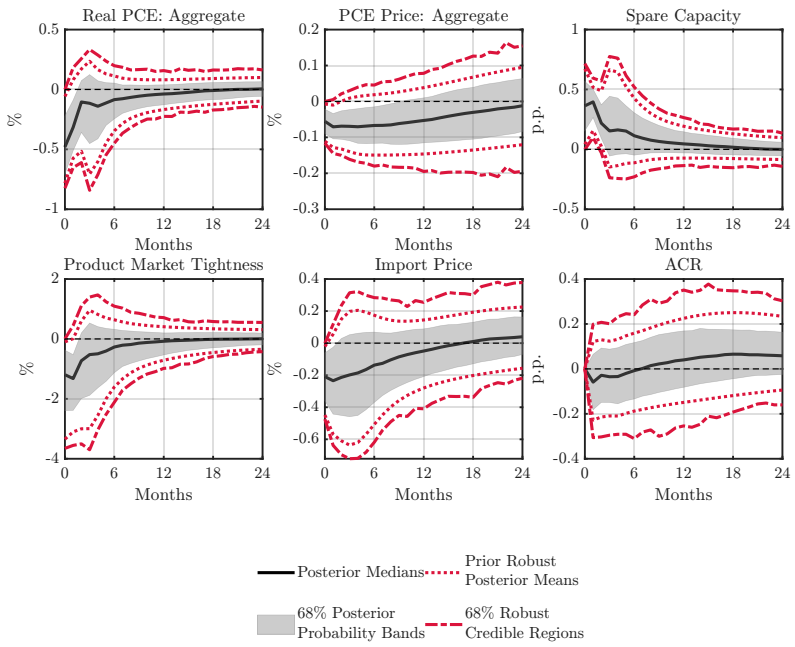


Figure G.1: IRFs to an Adverse Aggregate Demand Shock: Prior Robustness

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

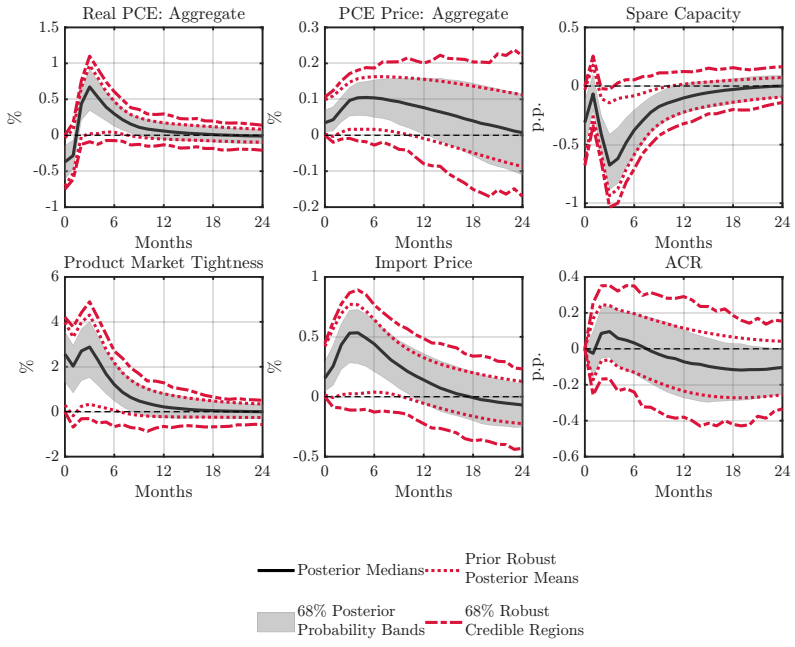


Figure G.2: IRFs to an Adverse Productive Capacity Shock: Prior Robustness

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

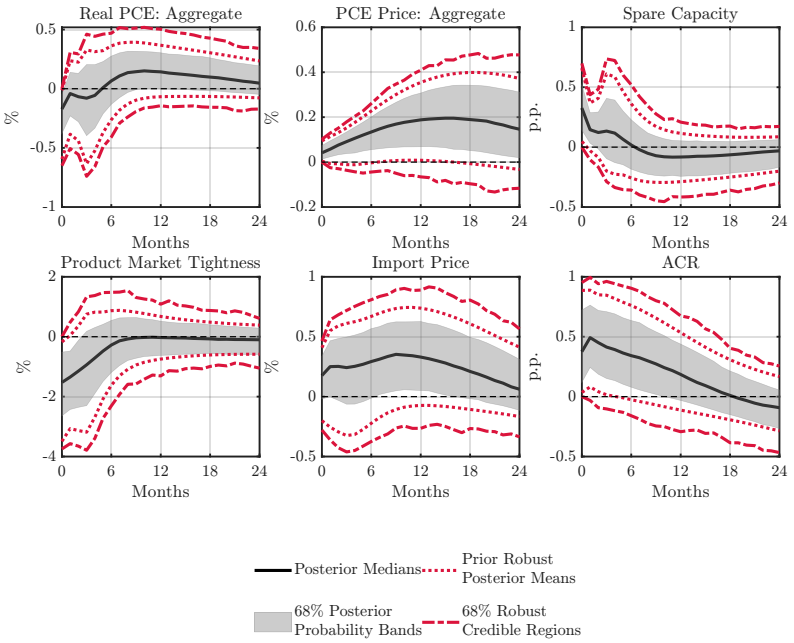


Figure G.3: IRFs to an Adverse Supply Chain Shock: Prior Robustness

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

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

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

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

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

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

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

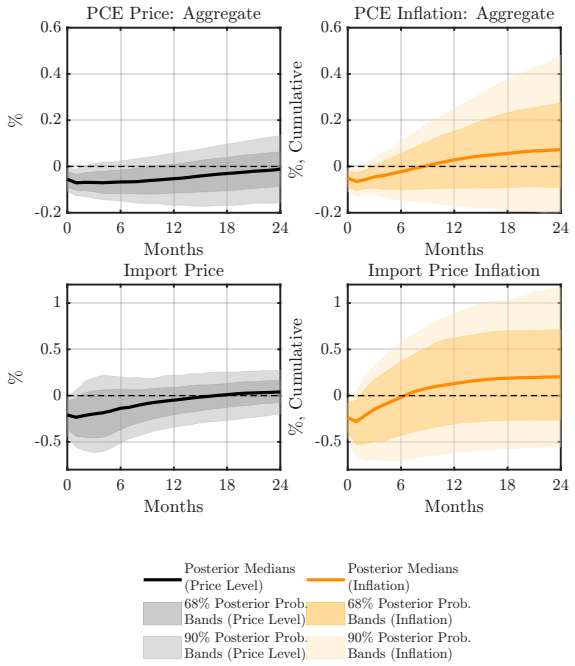


Figure G.4: Price-Level vs. Cumulated-Inflation IRFs: Adverse Aggregate Demand Shock

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

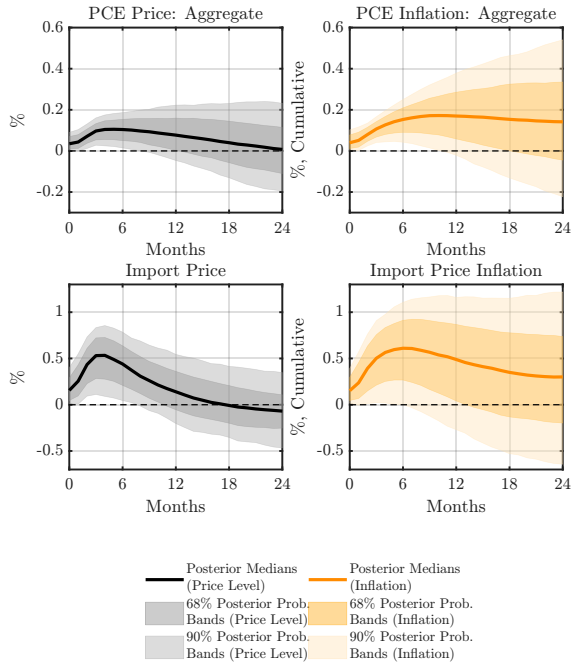


Figure G.5: Price-Level vs. Cumulated-Inflation IRFs: Adverse Productive Capacity Shock

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

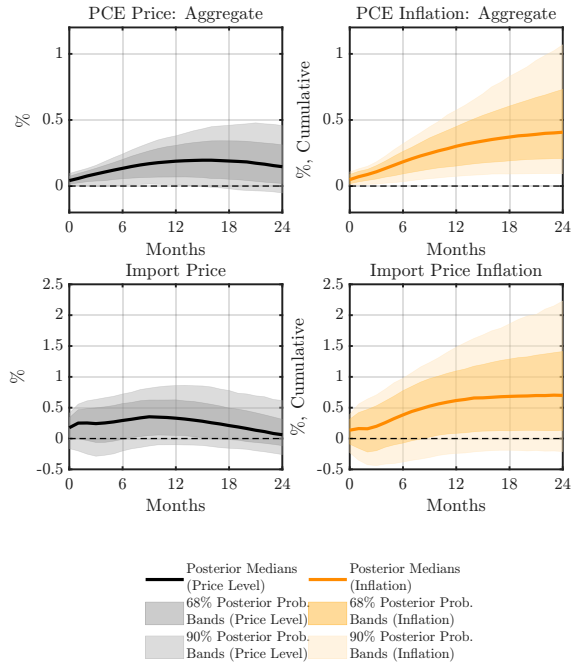


Figure G.6: Price-Level vs. Cumulated-Inflation IRFs: Adverse Supply Chain Shock

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

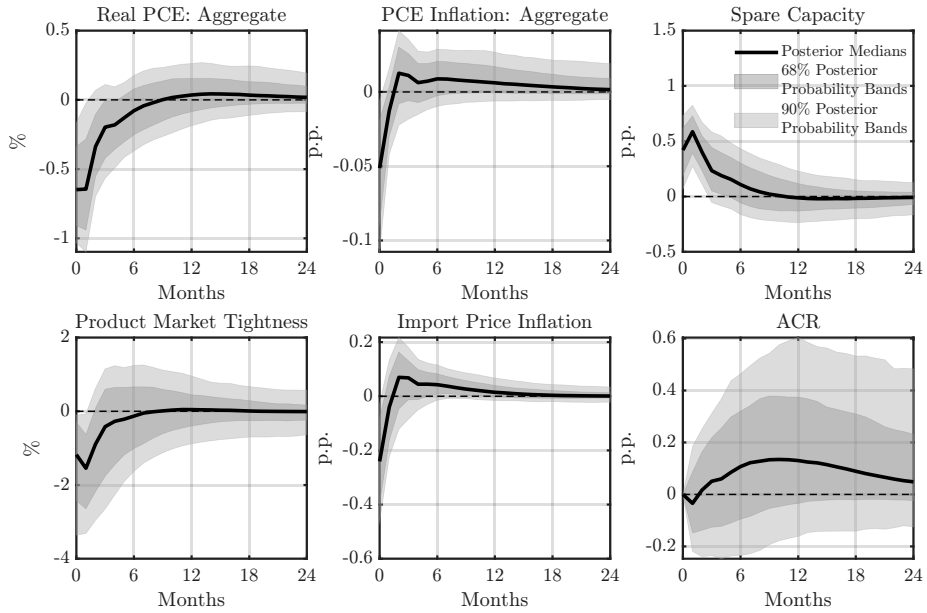


Figure G.7: IRFs to an Adverse Aggregate Demand Shock: PCE and Import Price Inflation

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

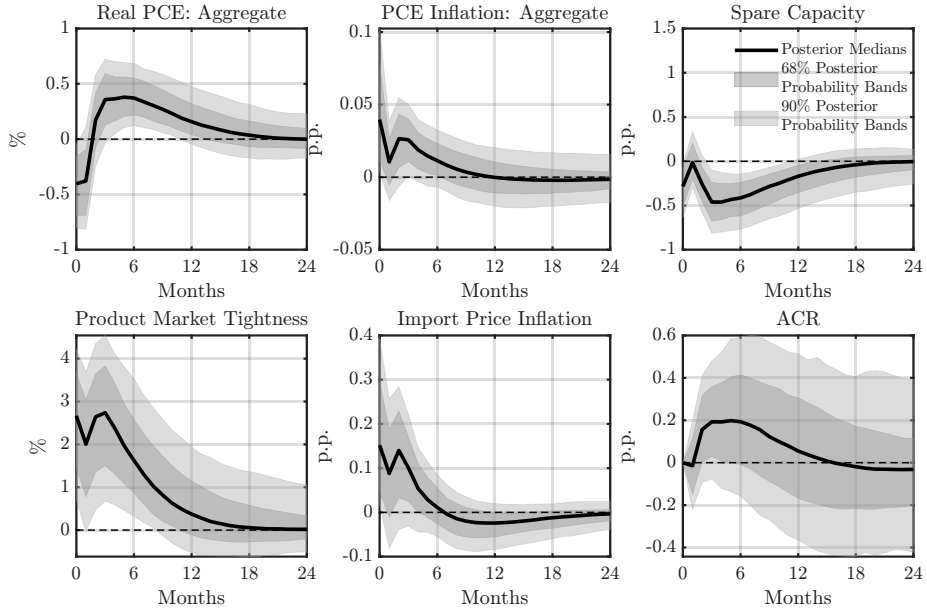


Figure G.8: IRFs to an Adverse Productive Capacity Shock: PCE and Import Price Inflation

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

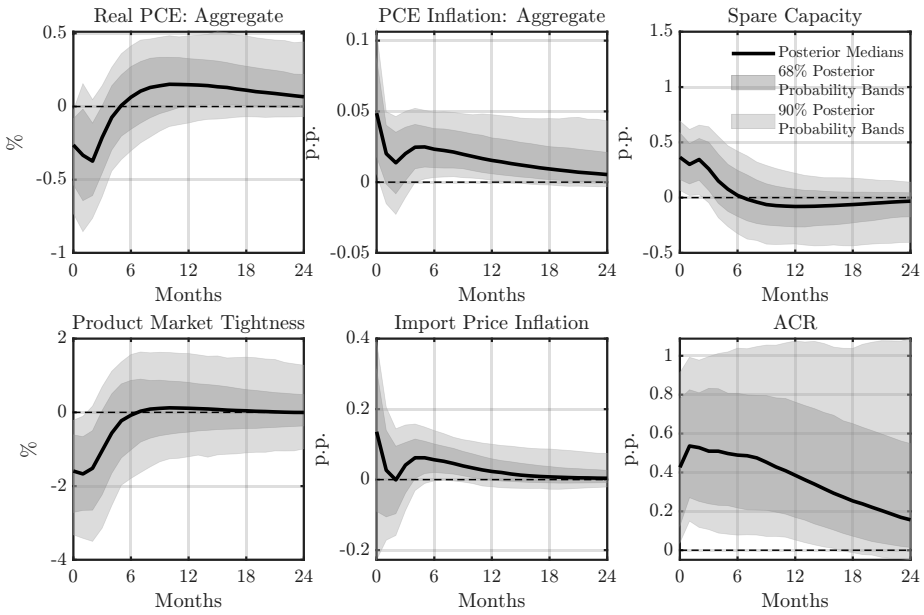


Figure G.9: IRFs to an Adverse Supply Chain Shock: PCE and Import Price Inflation

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

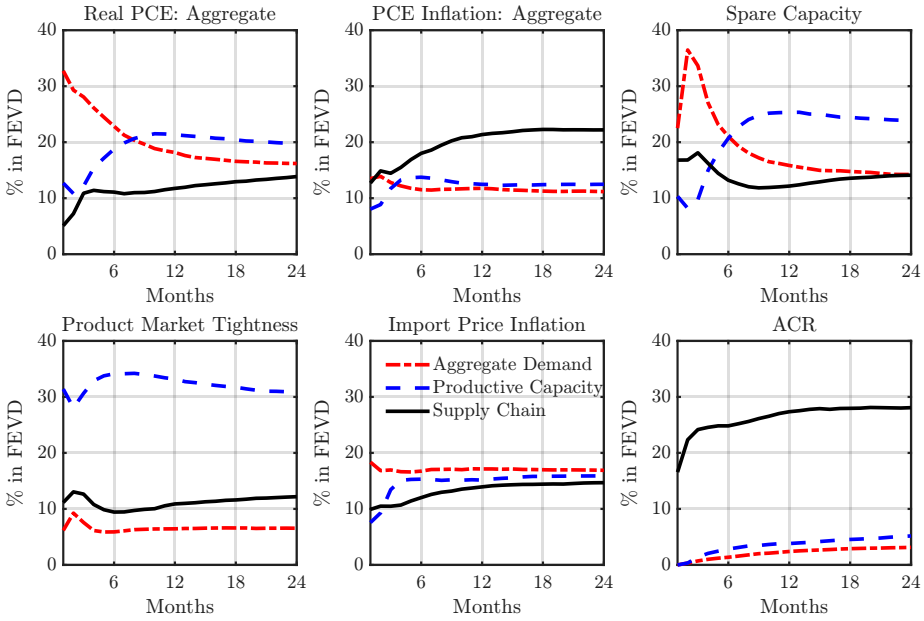


Figure G.10: FEVDs from the SVAR: PCE and Import Price Inflation

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

G.3. Invertibility Under an Expanded Information Set

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

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

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

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

[14] We thank an anonymous referee for suggesting this check.

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

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

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

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

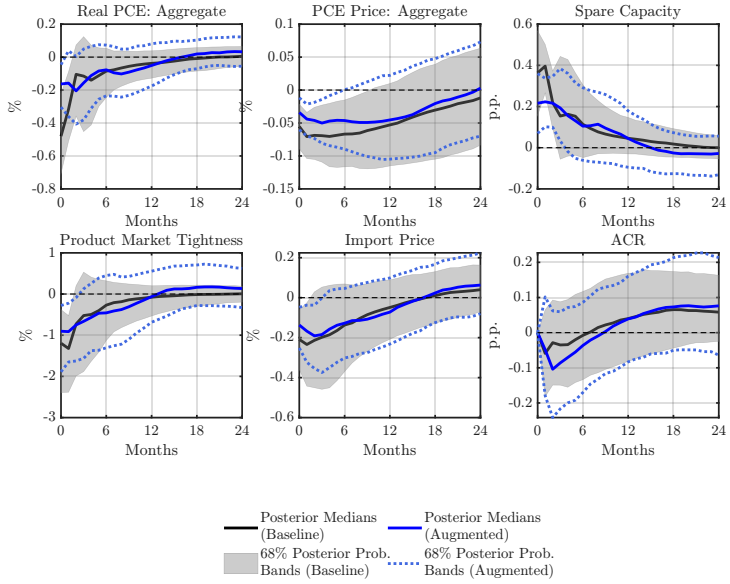


Figure G.11: IRFs to an Adverse Aggregate Demand Shock: Baseline vs. Augmented

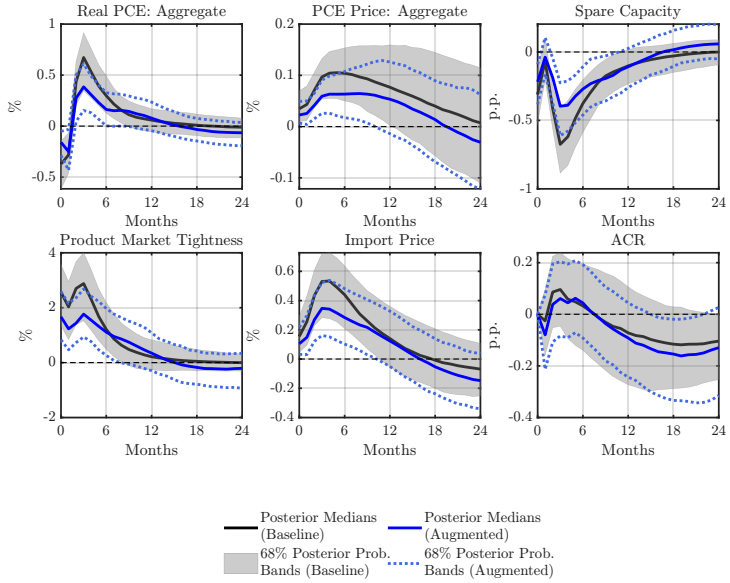


Figure G.12: IRFs to an Adverse Productive Capacity Shock: Baseline vs. Augmented

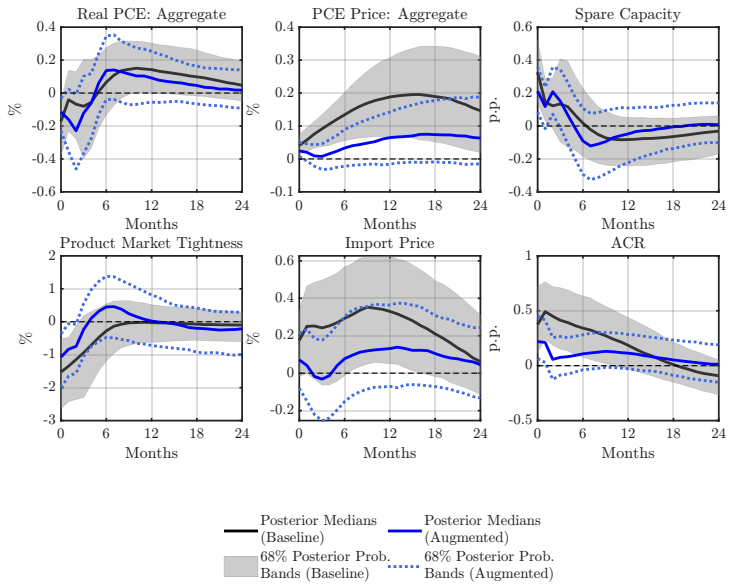


Figure G.13: IRFs to an Adverse Supply Chain Shock: Baseline vs. Augmented

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

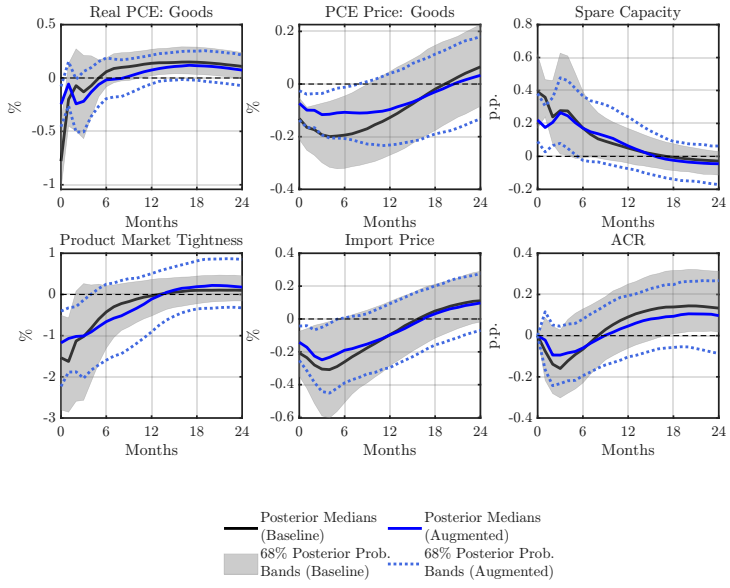


Figure G.14: IRFs to an Adverse Goods Demand Shock: Goods-Focused vs. Augmented

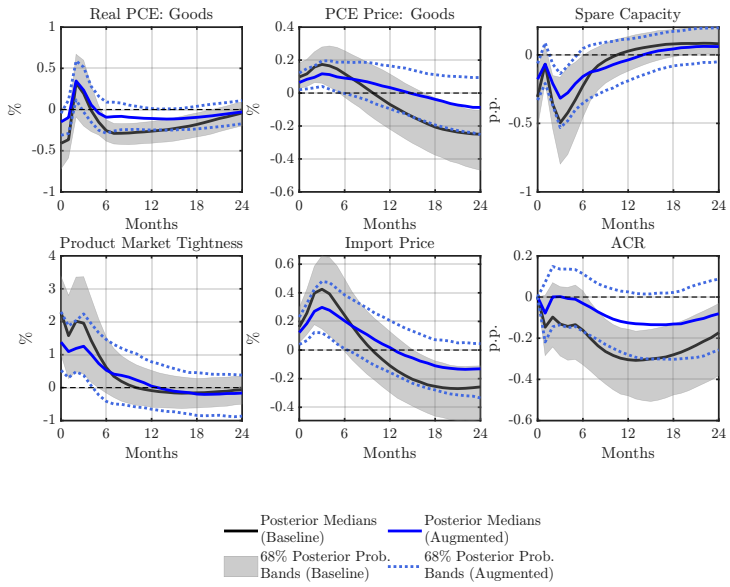


Figure G.15: IRFs to an Adverse Productive Capacity Shock: Goods-Focused vs. Augmented

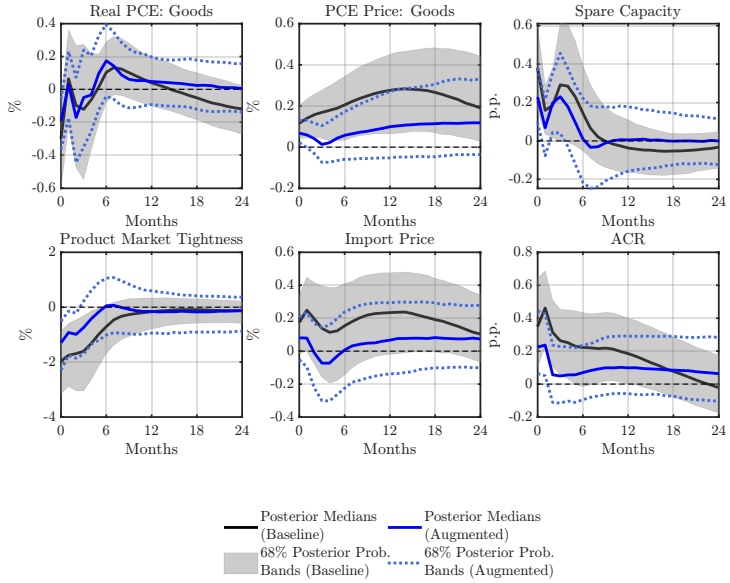


Figure G.16: IRFs to an Adverse Supply Chain Shock: Goods-Focused vs. Augmented

G.4. Nonlinear Transmission by Shock Size

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

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

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

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

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

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

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

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

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

The shock-size-specific IRFs are:

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

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

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

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

[17] See Appendix B for details on SLP.

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

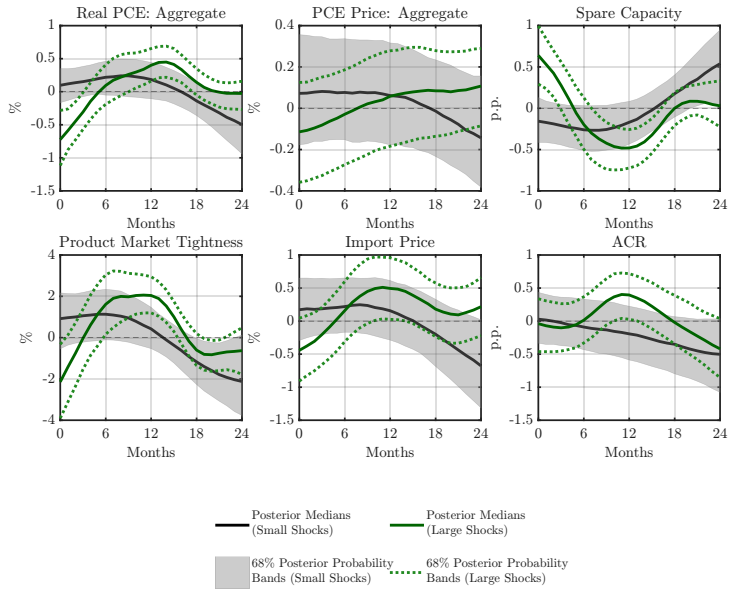


Figure G.17: Nonlinear Transmission of Aggregate Demand Shocks by Size

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

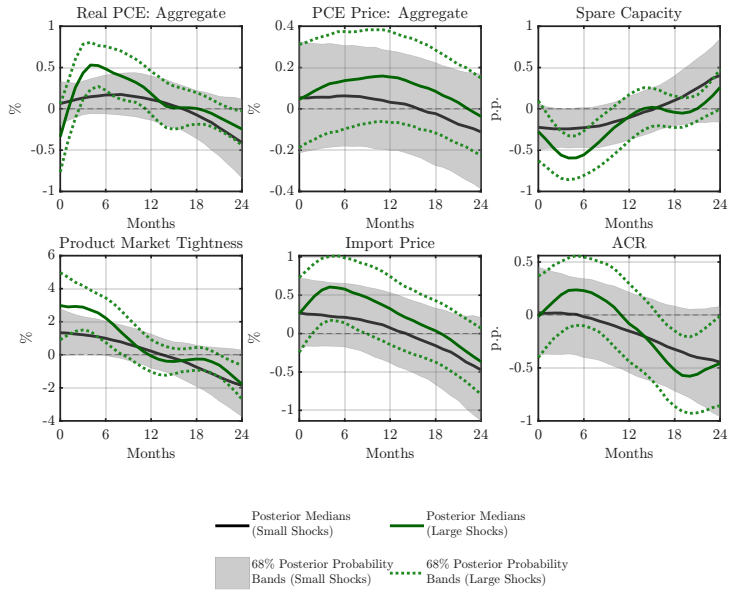


Figure G.18: Nonlinear Transmission of Productive Capacity Shocks by Size

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

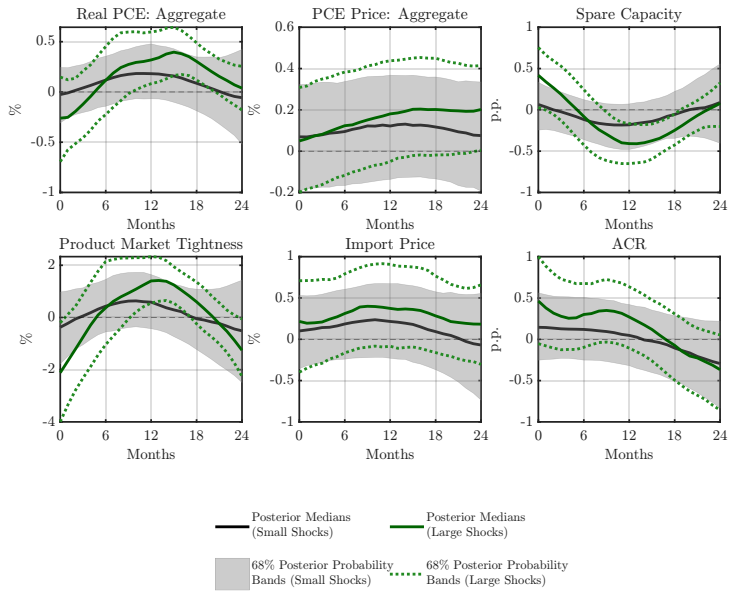


Figure G.19: Nonlinear Transmission of Supply Chain Shocks by Size

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

G.5. Different Lag Structures

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

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

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

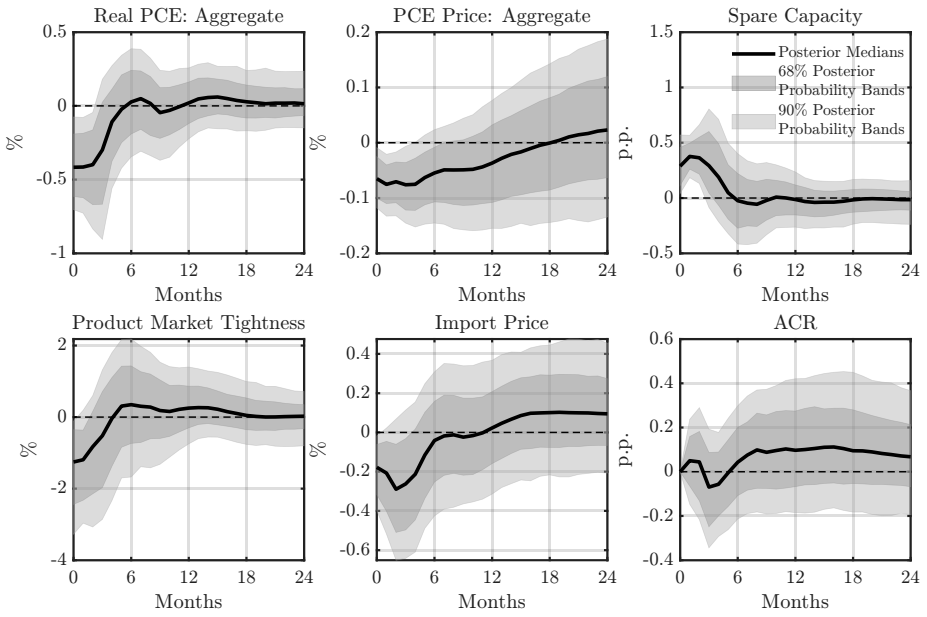


Figure G.20: IRFs to an Adverse Aggregate Demand Shock: Four Lags

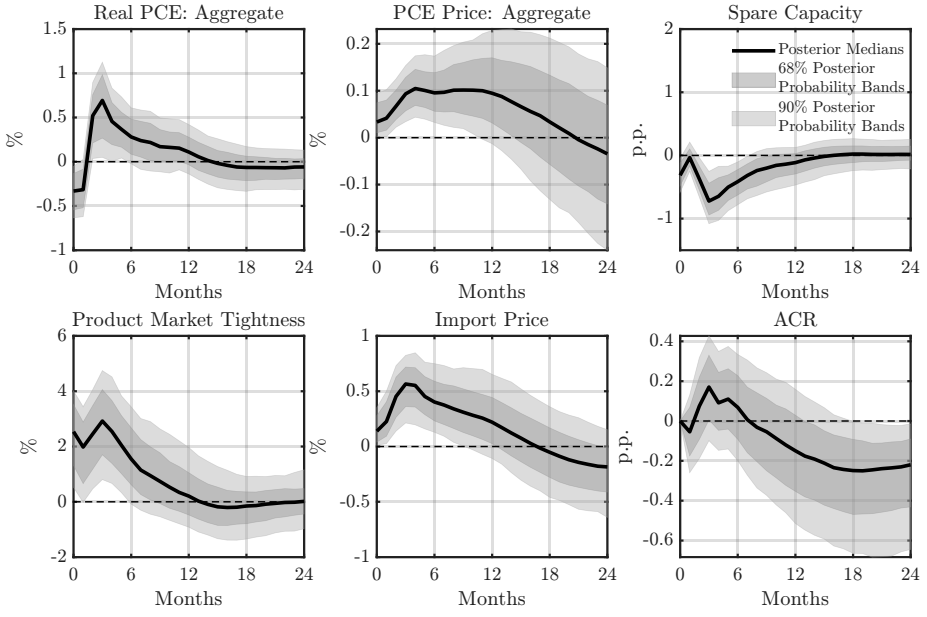


Figure G.21: IRFs to an Adverse Productive Capacity Shock: Four Lags

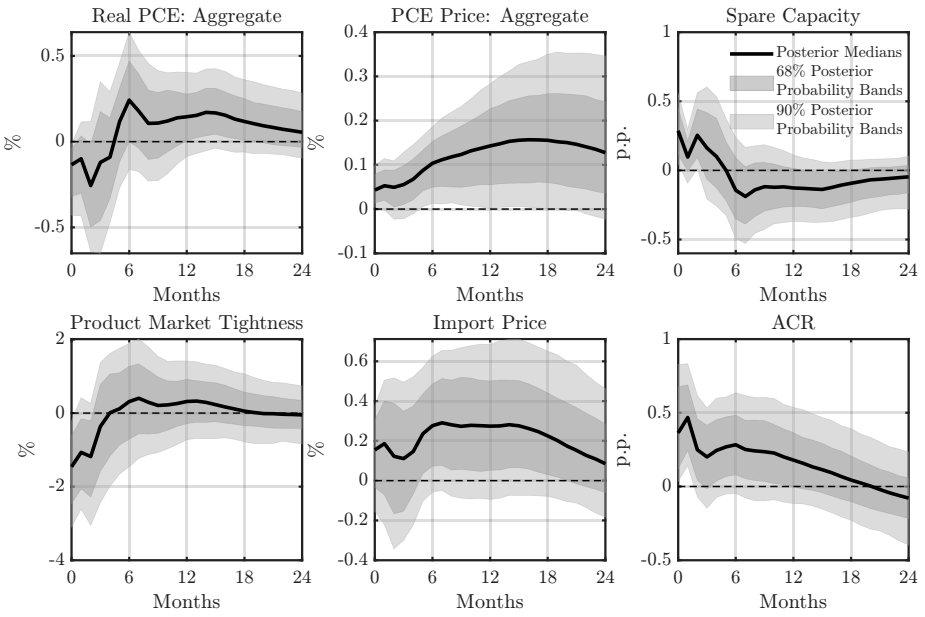


Figure G.22: IRFs to an Adverse Supply Chain Shock: Four Lags

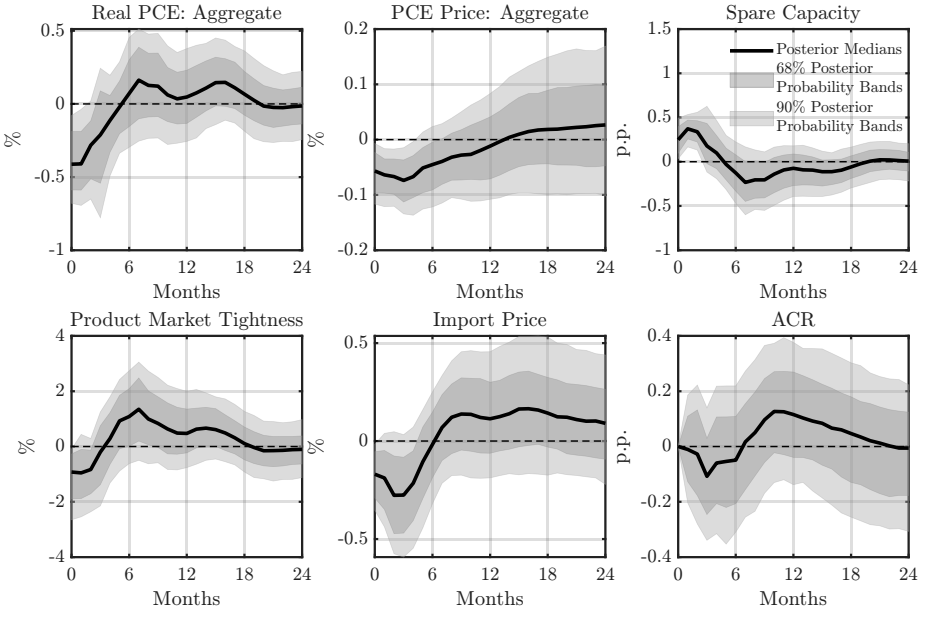


Figure G.23: IRFs to an Adverse Aggregate Demand Shock: Six Lags

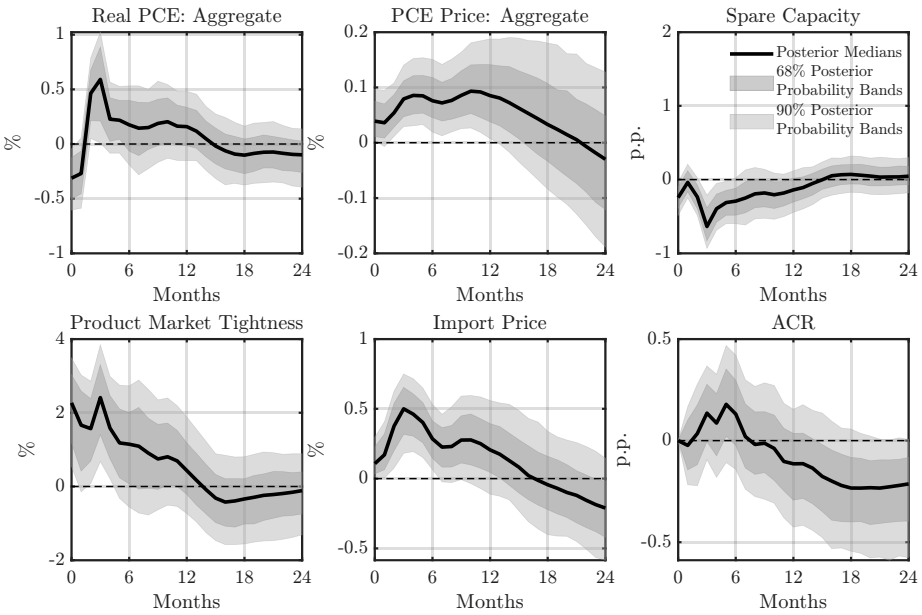


Figure G.24: IRFs to an Adverse Productive Capacity Shock: Six Lags

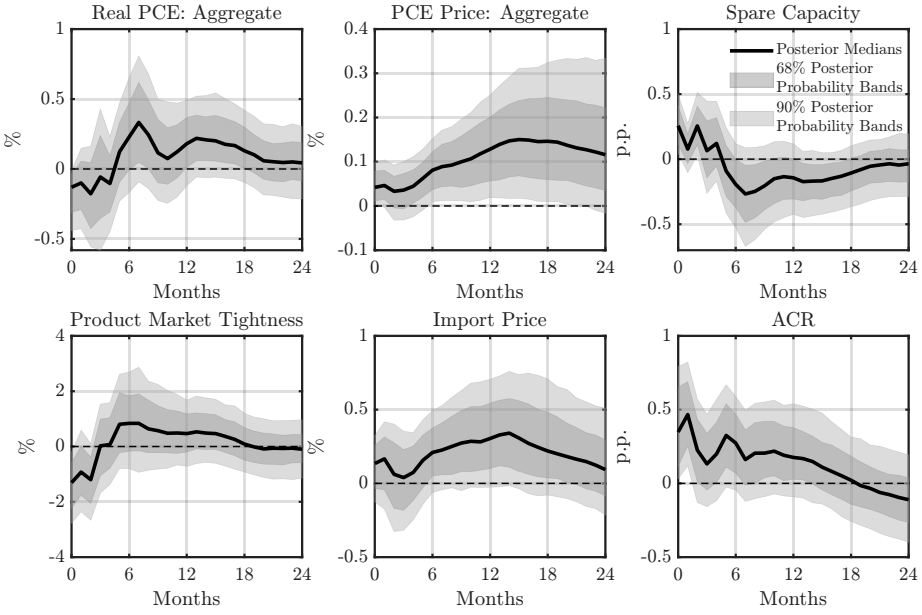


Figure G.25: IRFs to an Adverse Supply Chain Shock: Six Lags

G.6. Dropping the Linear Trend

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

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

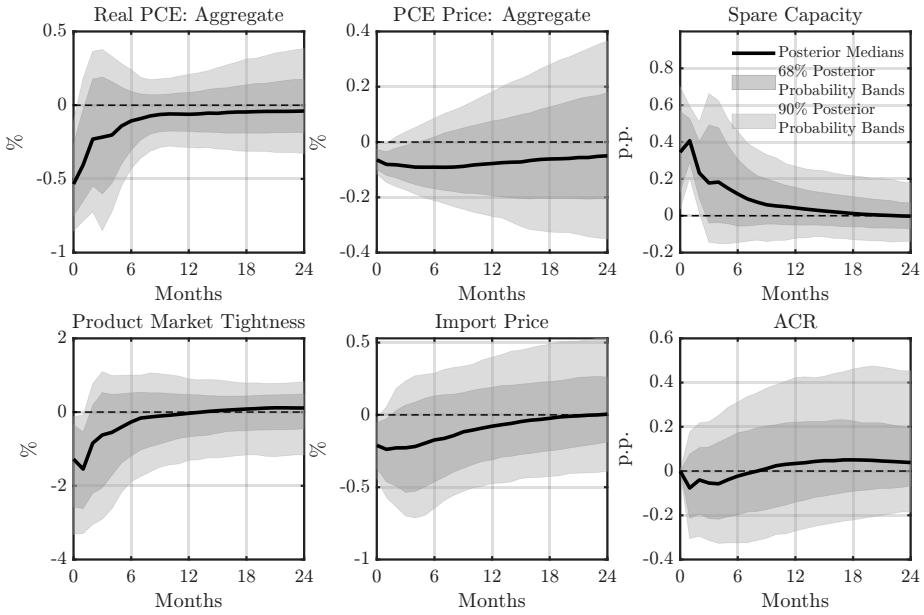


Figure G.26: IRFs to an Adverse Aggregate Demand Shock: Constant Only

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

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

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

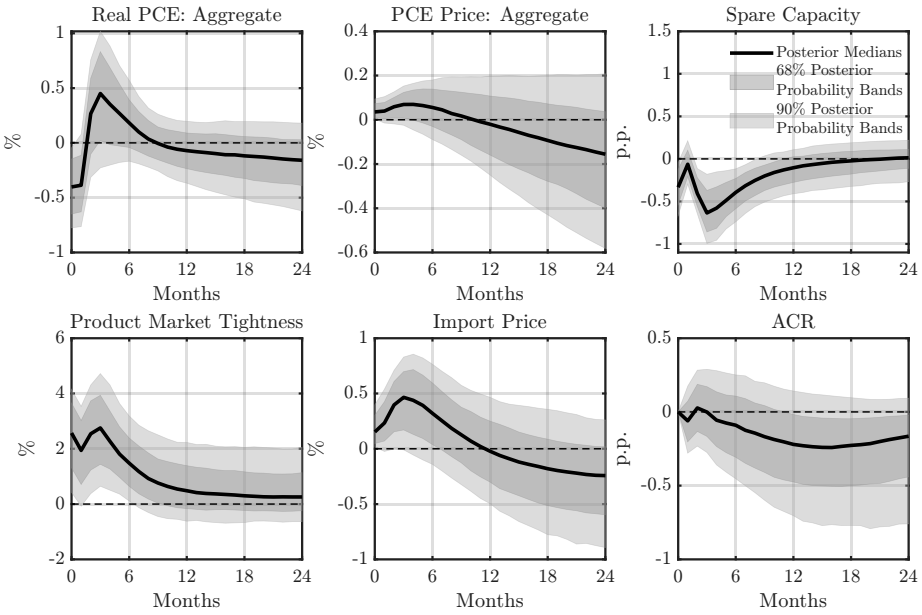


Figure G.27: IRFs to an Adverse Productive Capacity Shock: Constant Only

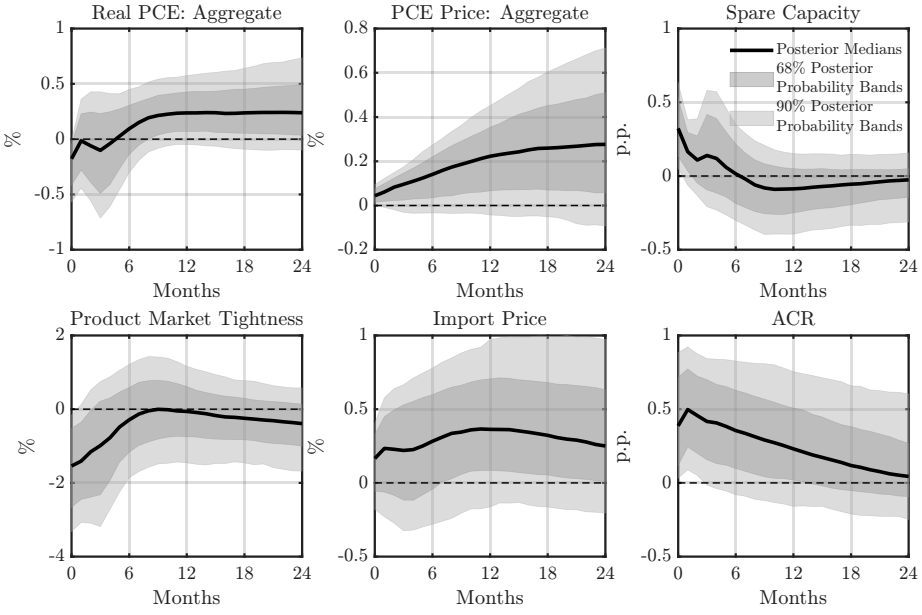


Figure G.28: IRFs to an Adverse Supply Chain Shock: Constant Only

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

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

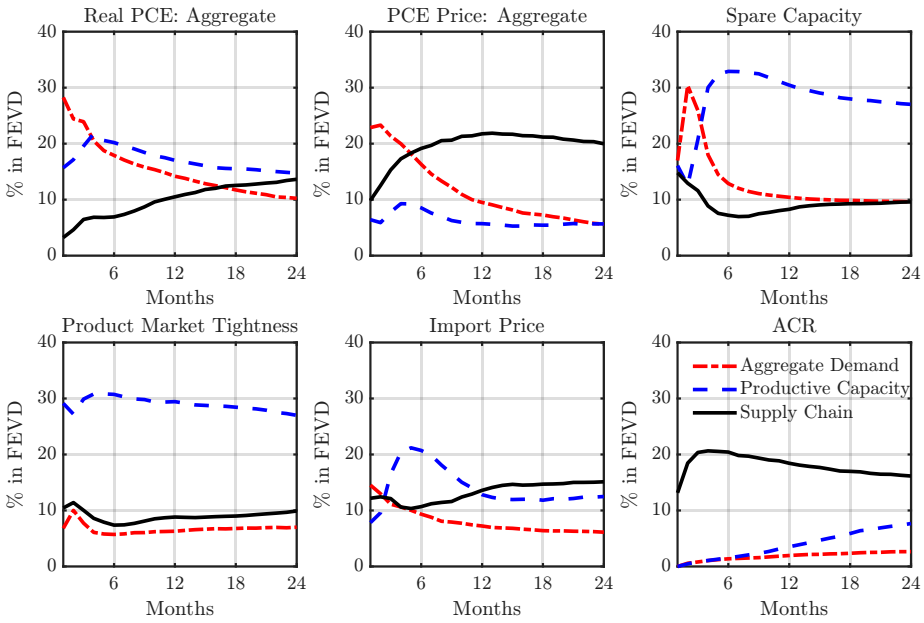


Figure G.29: FEVDs from the SVAR: Constant Only

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

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

G.7. Alternative Proxies for Activity and Prices

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

following Chow and Lin (1971). All series are seasonally adjusted.

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

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

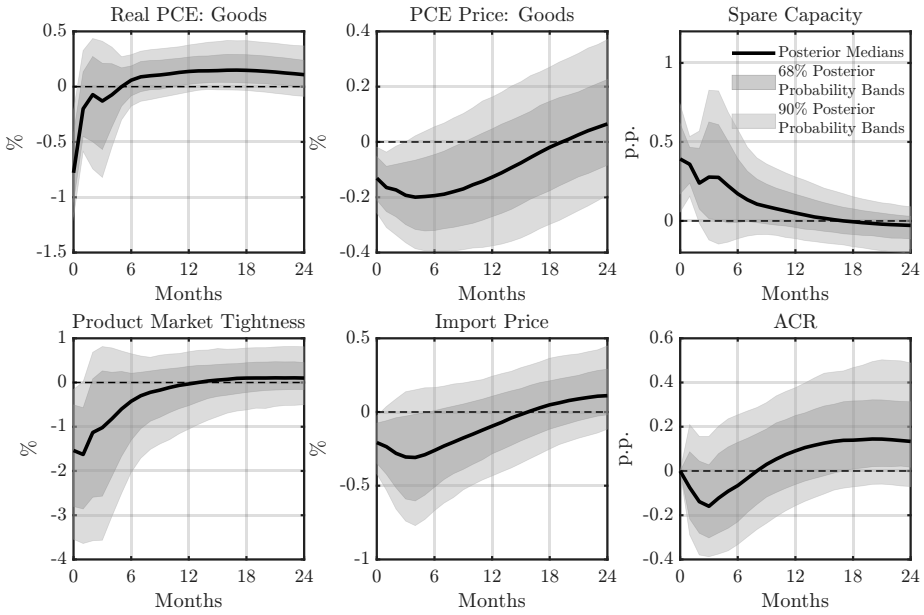


Figure G.30: IRFs to an Adverse Goods Demand Shock: Goods PCE and Prices

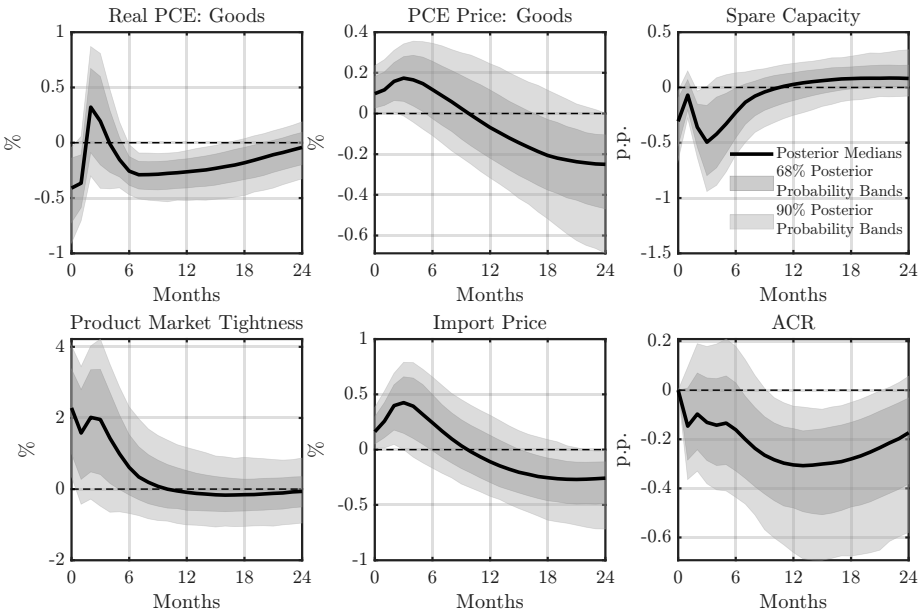


Figure G.31: IRFs to an Adverse Productive Capacity Shock: Goods PCE and Prices

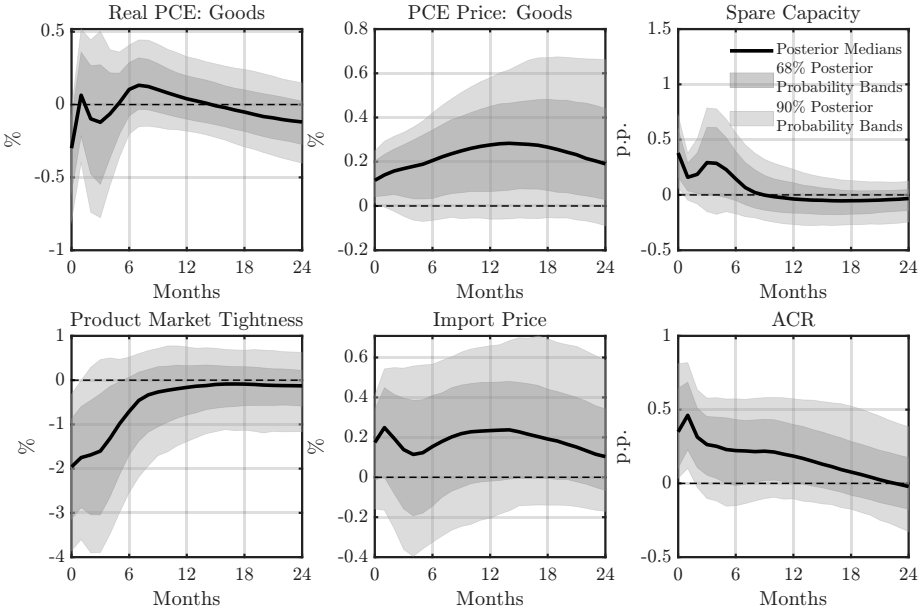


Figure G.32: IRFs to an Adverse Supply Chain Shock: Goods PCE and Prices

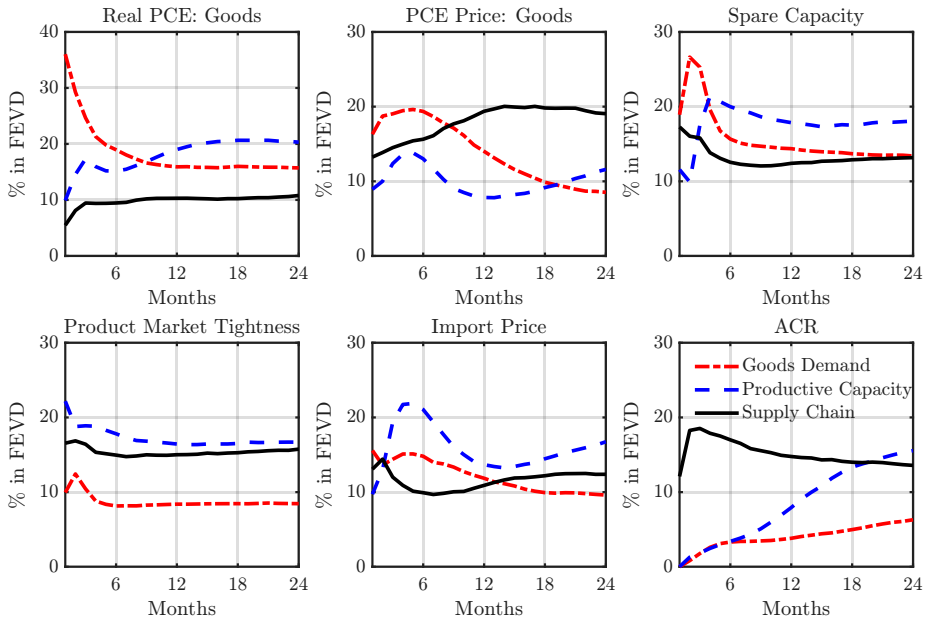


Figure G.33: FEVDs from the SVAR: Goods PCE and Prices

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

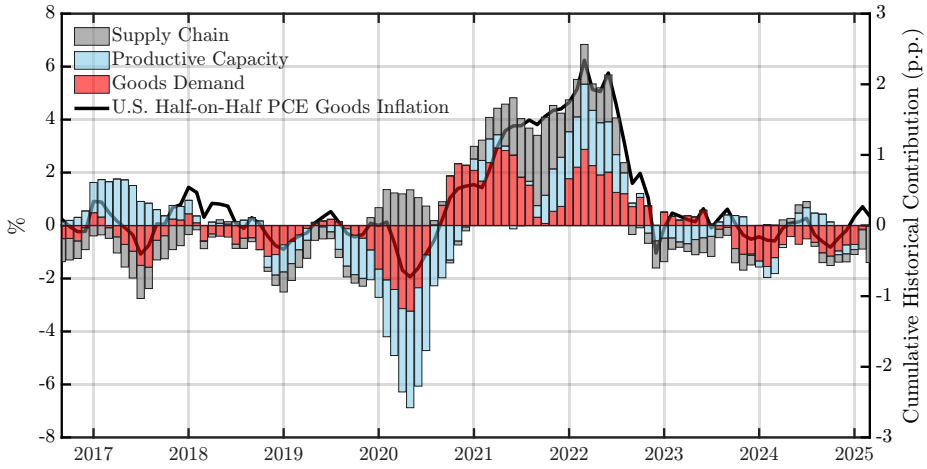


Figure G.34: HD of U.S. HoH PCE Goods Inflation

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

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

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

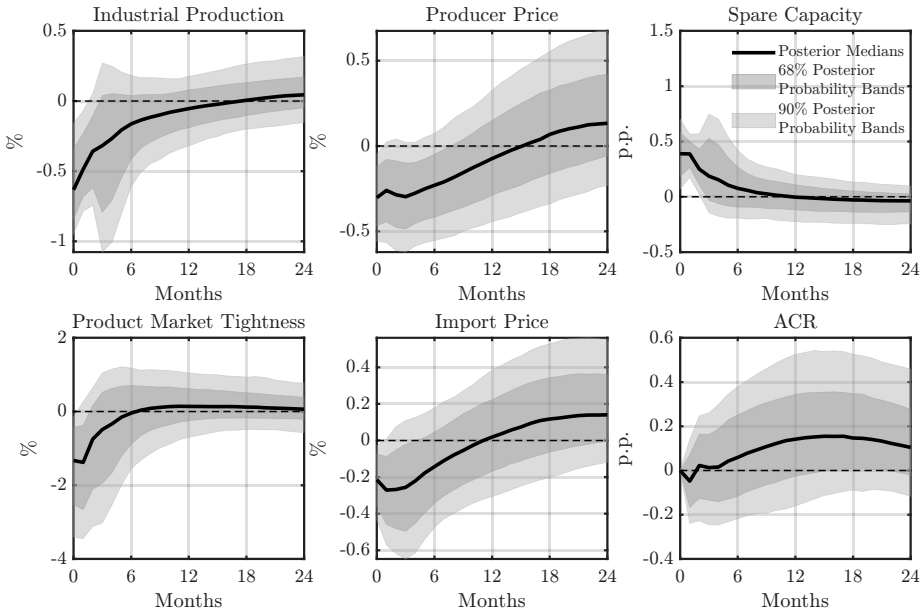


Figure G.35: IRFs to an Adverse Goods Demand Shock: IP and PPI

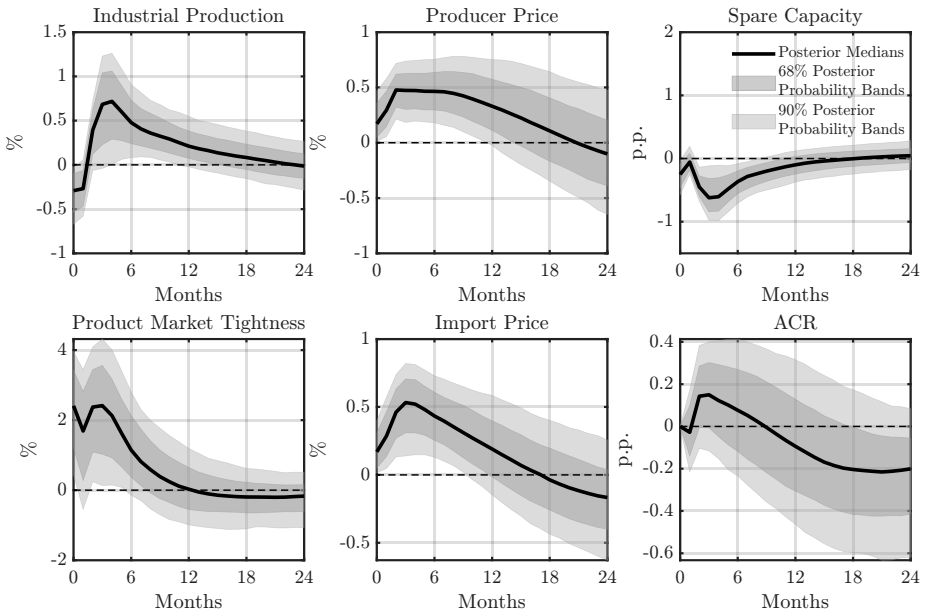


Figure G.36: IRFs to an Adverse Productive Capacity Shock: IP and PPI

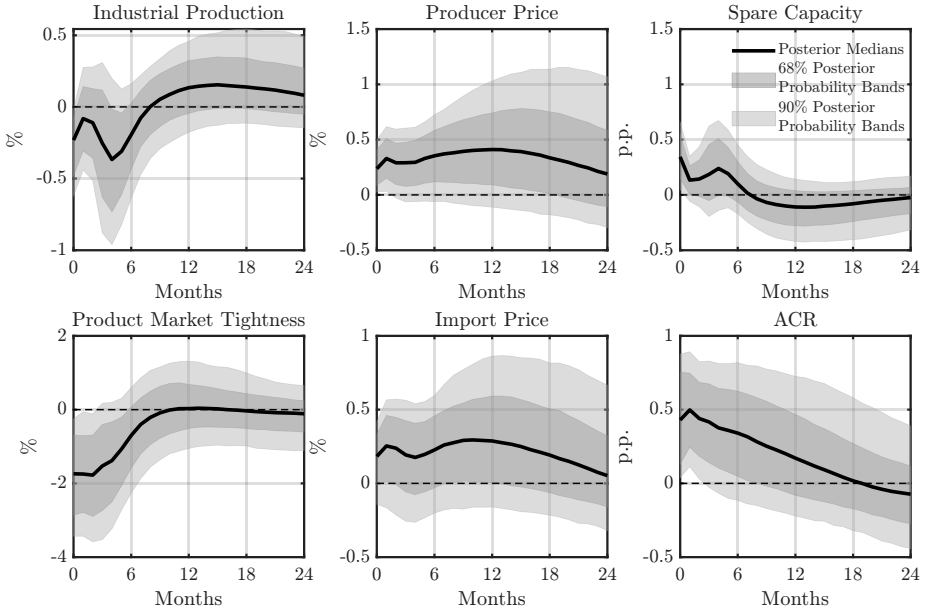


Figure G.37: IRFs to an Adverse Supply Chain Shock: IP and PPI

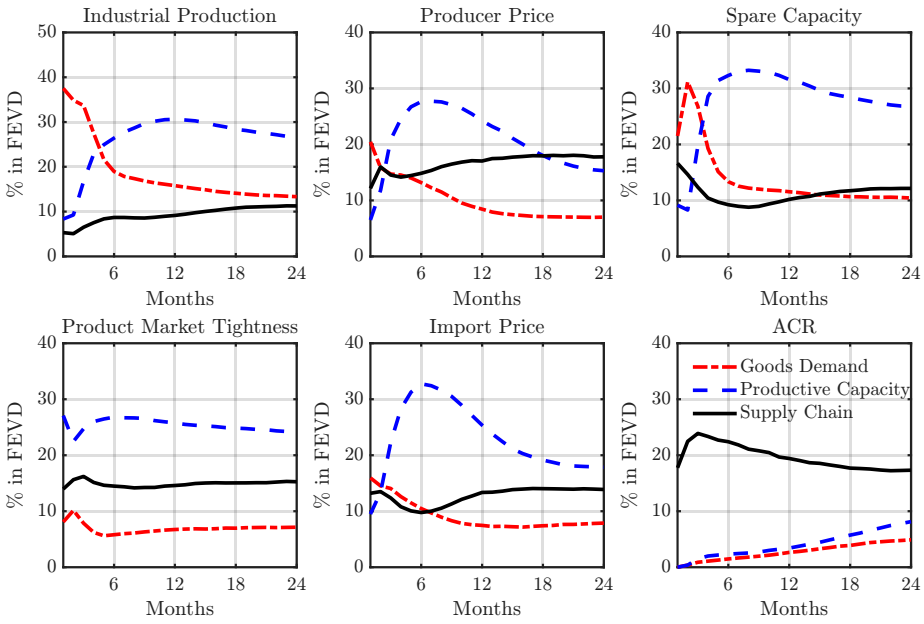


Figure G.38: FEVDs from the SVAR: IP and PPI

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

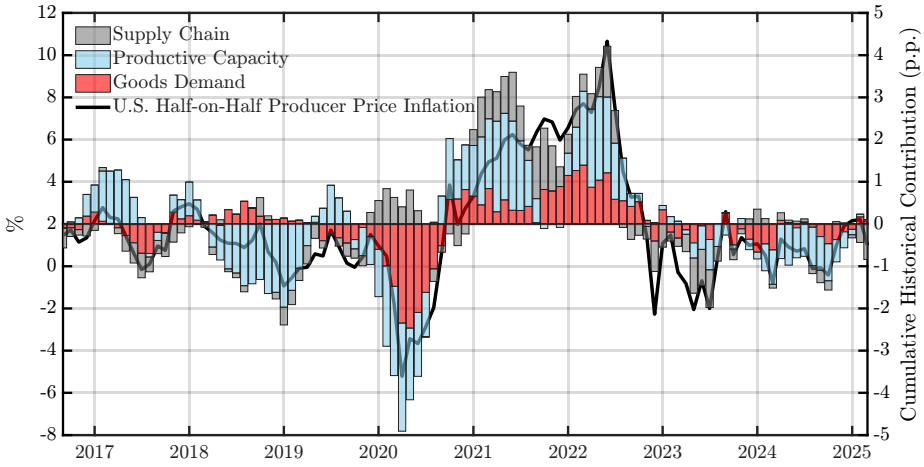


Figure G.39: HD of U.S. HoH Producer Price Inflation

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

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

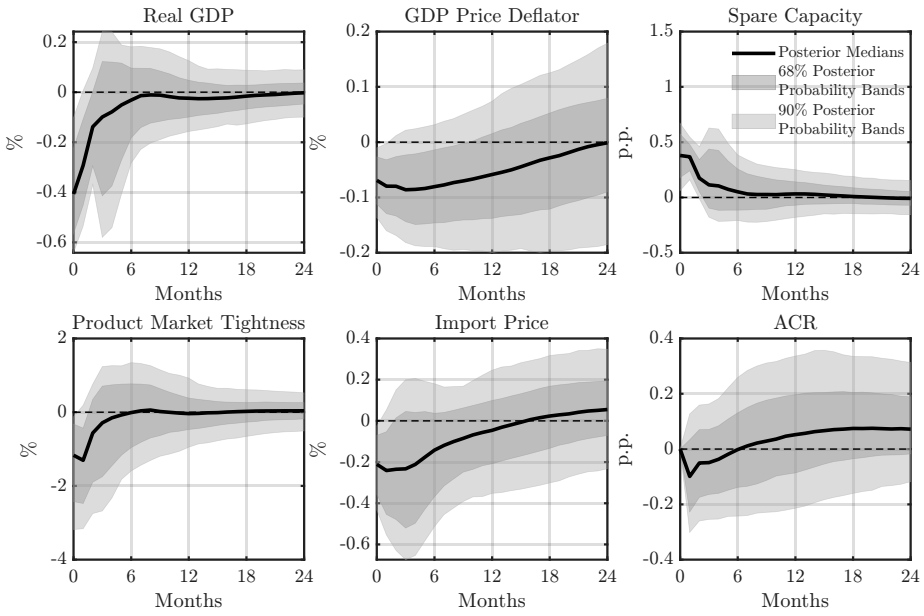


Figure G.40: IRFs to an Adverse Aggregate Demand Shock: Real GDP and GDP Deflator

Taking stock. In sum, four observations stand out. First, the macroeconomic impacts of demand and capacity shocks are robust across alternative measures of output and prices. Second, supply chain price effects are at least as strong, and often stronger, when producer- or GDP-based price indices are used, underscoring the role of logistics bottlenecks as an independent source of price variation. Third, the FEVDs consistently show that (i) capacity shocks dominate real activity at medium horizons, (ii) supply chain shocks account for a sizable share of price variance, and (iii) demand shocks remain concentrated at short horizons. This pattern holds in both the baseline FEVDs and all three alternatives. Finally, across different proxies for inflation, a large

share of the post-pandemic surge in U.S. inflation can be attributed to supply-side disruptions, with the relative importance of capacity versus supply chain shocks varying by measure.

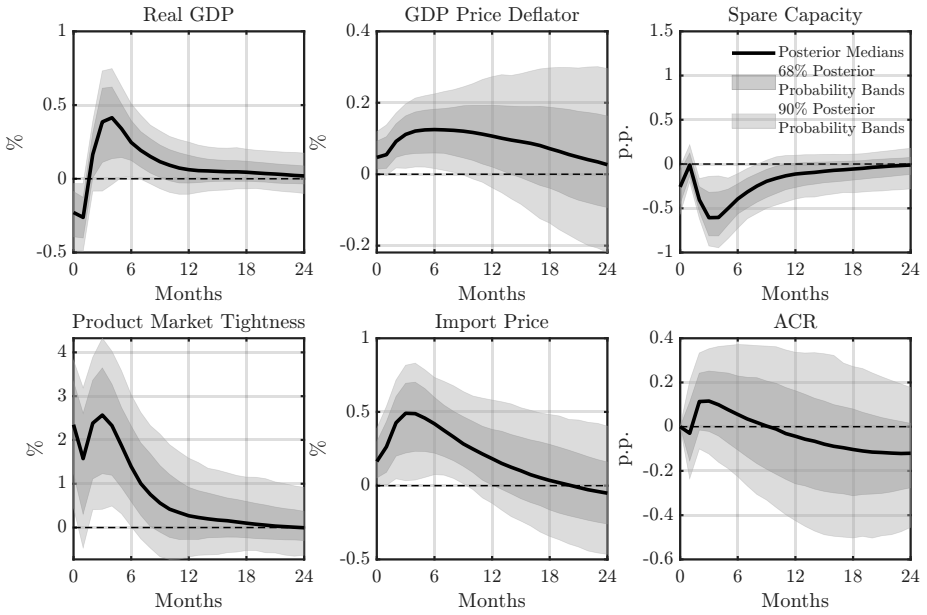


Figure G.41: IRFs to an Adverse Productive Capacity Shock: Real GDP and GDP Deflator

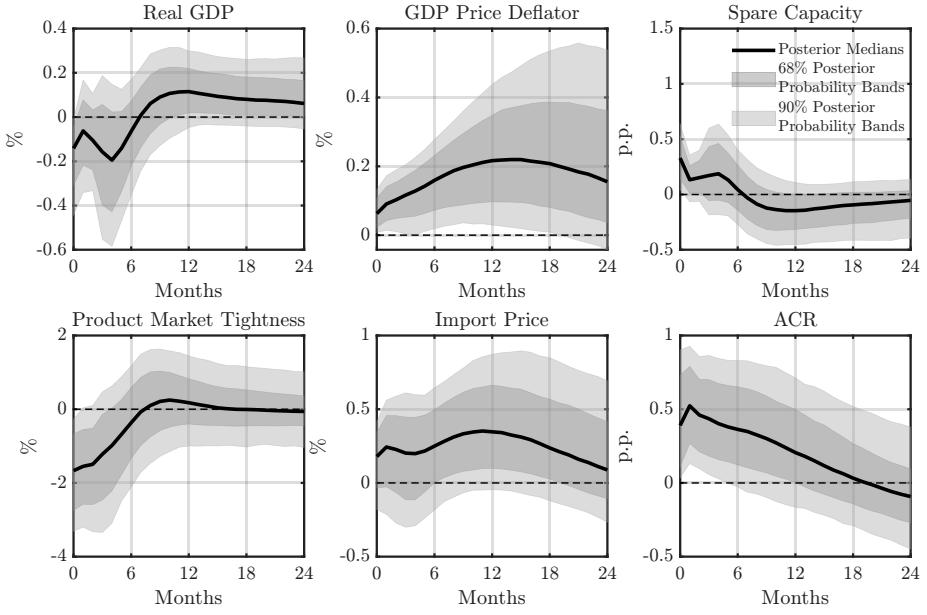


Figure G.42: IRFs to an Adverse Supply Chain Shock: Real GDP and GDP Deflator

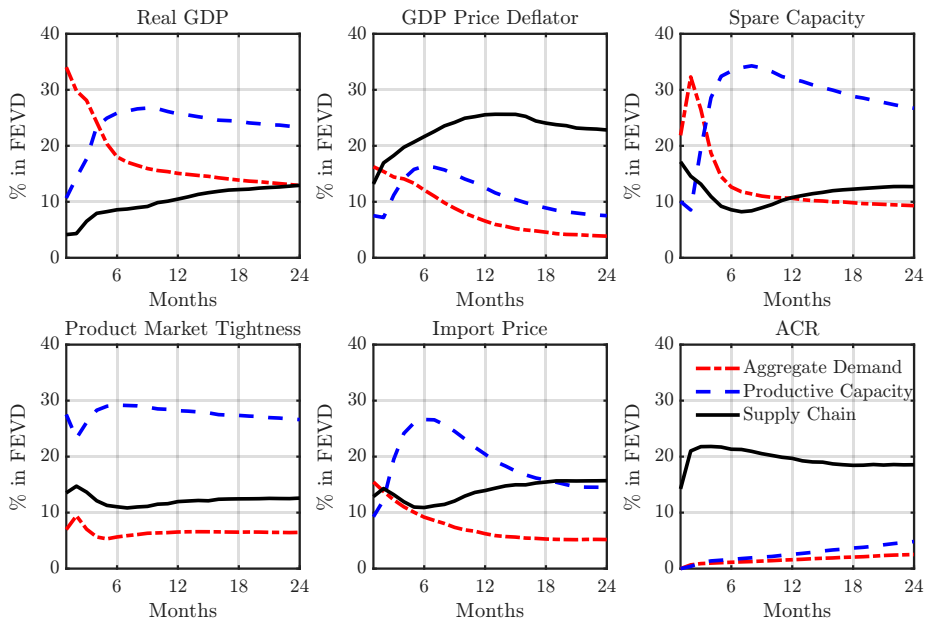


Figure G.43: FEVDs from the SVAR: Real GDP and GDP Deflator

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

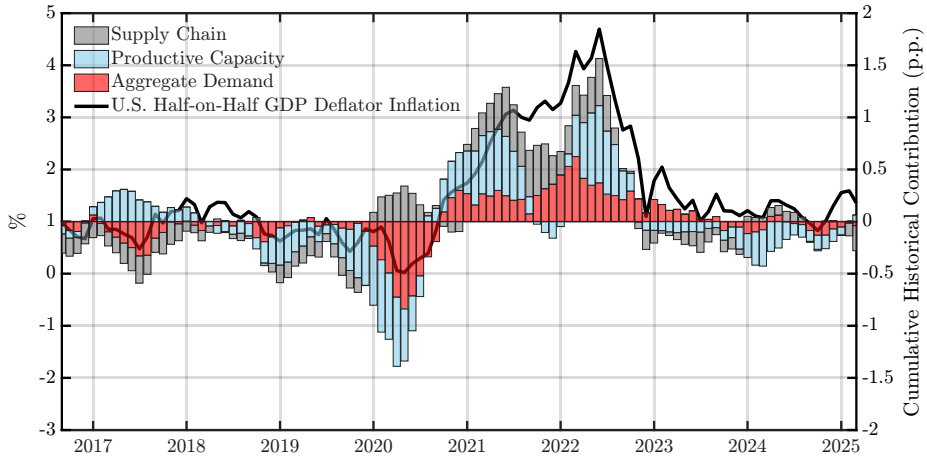


Figure G.44: HD of U.S. HoH GDP Deflator Inflation

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

H. Identification Gains From Supply Chain Proxies

This appendix documents the identification gains —both in the magnitude of price responses to supply chain disturbances and in the precision of the posterior estimates— that arise when our congestion indices are used as measures of global supply chain conditions in the SVAR.

As a starting point, Appendix H.1 estimates the causal effects of supply chain shocks when no proxy is included, providing a natural lower bound against which to evaluate the gains from adding proxies and imposing domain-knowledge-based identification restrictions. Appendix H.2 then includes the ACR index but removes the zero restrictions on its responses to aggregate demand and productive capacity shocks at horizon $k = 1$, showing that much of the improvement comes simply from incorporating our congestion indices.

Appendices H.3 and H.4 show that using the ACT index or a targeted ACR index for major Trans-Pacific ports central to the U.S. goods supply chain delivers results that are quantitatively similar to those from the ACR baseline. By contrast, Appendices H.5-H.7 examine HARPEX, the New York Fed's GSCPI (Benigno et al., 2022), and the SDI from Smirnyagin and Tsyvinski (2022) and Liu, Smirnyagin and Tsyvinski (2024) as alternative proxies, revealing sizable discrepancies that materially alter the estimated macroeconomic effects of supply chain disturbances.

Finally, Appendix H.8 replaces the PCE chain-type price index with a goods-only measure, confirming that the main conclusions remain intact.

H.1. No Proxy

We re-estimate the SVAR without a direct proxy for global supply chain conditions and compare it with the baseline specification that includes ACR under zero restrictions on its responses to aggregate demand and productive capacity shocks at horizon $k = 1$ (Restrictions 1-3). The price IRFs in Figure H.1 clearly illustrate the loss of identification: without a proxy, the median responses of the PCE chain-type price index and the import price index to an adverse supply chain shock are smaller, and the 68% and 90% posterior bands overlap zero across most horizons. By contrast, with ACR and zero restrictions (Figure 9 in the main text), the median responses are larger, the bands are tighter, and they remain above zero for longer horizons, yielding a sharper inflationary signal from supply chain disturbances.

The FEVDs in Figure H.2 convey the same pattern. In the no-proxy case, a larger share of the forecast error variance of PCE prices is attributed to aggregate demand, with only a modest

role for the supply chain shock. In the baseline with ACR and zero restrictions (Figure 10), the supply chain contribution rises at medium horizons, which shows that ACR isolates the relevant supply side component that would otherwise be absorbed by aggregate demand.

The HD of U.S. headline PCE inflation in Figure H.3 reinforces this point. Without a proxy, the 2021-22 inflation surge is attributed mainly to demand and capacity shocks, and the supply chain shock plays a limited role. In the baseline (Figure 11), the supply chain contribution grows visibly larger over the same period, consistent with congestion in maritime logistics and showing that ACR provides cleaner identification of supply-chain-driven price pressures.

Taken together, the no-proxy specification serves as a natural lower bound. Introducing ACR and imposing the $k = 1$ zero restrictions, motivated by industrial practice, increases the magnitude and precision of price responses and shifts the variance and historical contributions toward the supply chain shock, in line with the mechanism emphasized in our framework.

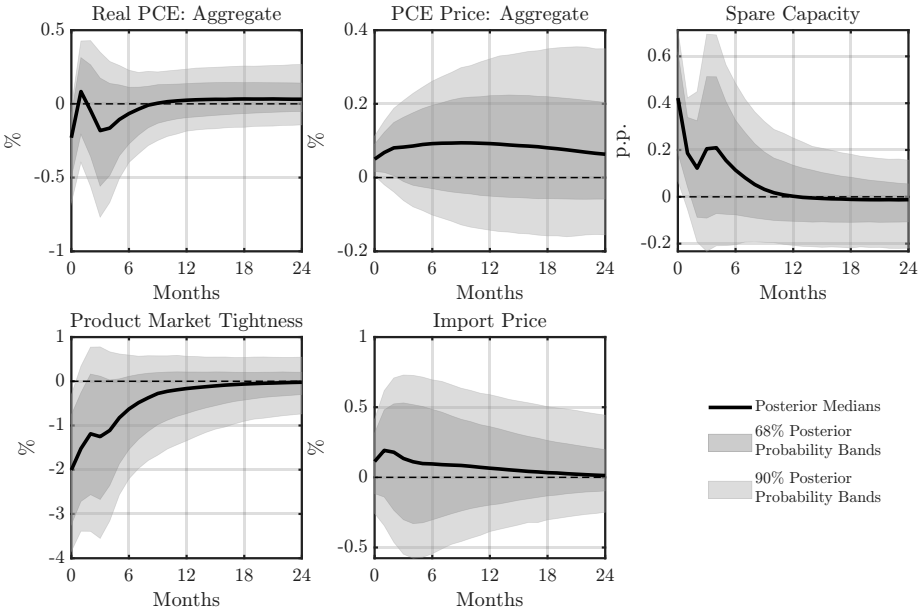


Figure H.1: IRFs to an Adverse Supply Chain Shock: No Proxy

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification described in Section III, except that the ACR index is omitted.

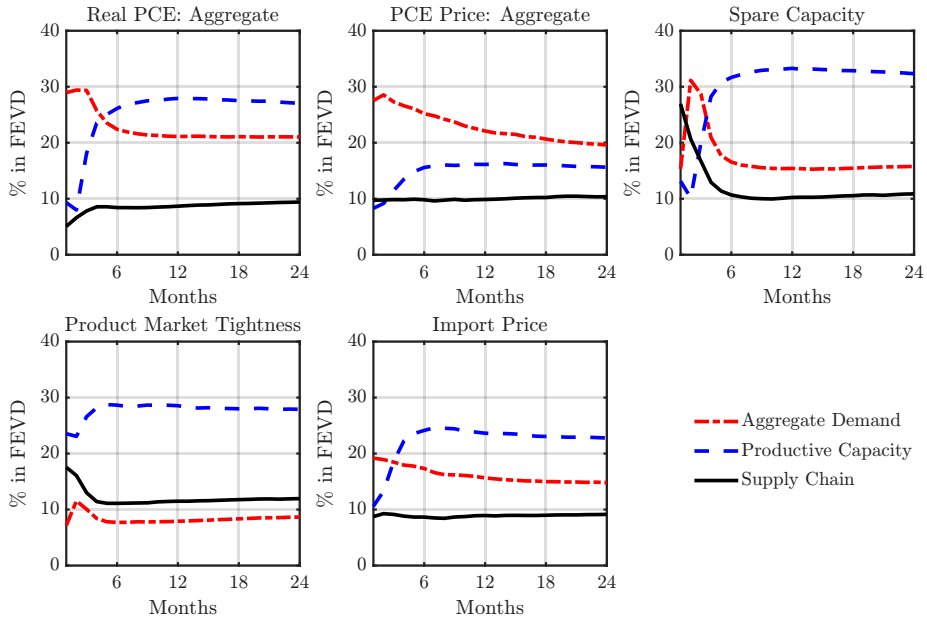


Figure H.2: FEVDs from the SVAR: No Proxy

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the ACR index is omitted.

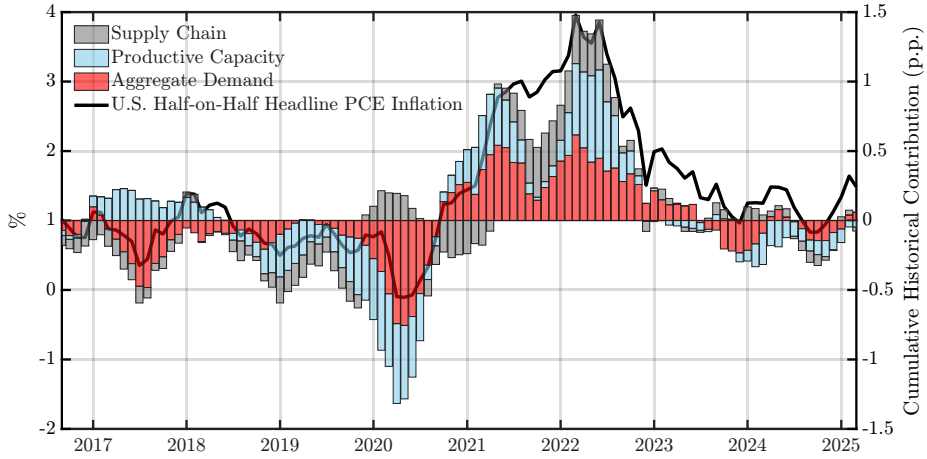


Figure H.3: HD of U.S. HoH Headline PCE Inflation: No Proxy

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR identical to the baseline, except that the ACR index is omitted.

H.2. ACR Without Zero Restrictions

Next, we include the ACR index in the SVAR but remove the zero restrictions on its responses to aggregate demand and productive capacity shocks at horizon $k = 1$. Figure H.4 shows that, relative to the no-proxy case, adding ACR raises the median response of the PCE chain-type price index to an adverse supply chain shock, with narrower 68% and 90% posterior bands that more often lie above zero. However, the responses remain less precisely estimated than in the baseline with ACR and zero restrictions (Figure 9 in the main text), where the bands tighten further and the inflationary signal is clearer.

Figures H.5 and H.6 point to the same conclusion. With the ACR index, the supply chain shock explains a larger share of the forecast error variance of PCE prices and contributes more to the 2021–22 inflation surge than in the no-proxy case. Yet, in both dimensions, its role remains smaller than in the baseline (Figures 10 and 11), where the timing and magnitude more closely align with observed congestion in maritime logistics.

In sum, including ACR alone accounts for most of the identification gains (both in the size of price responses and the precision of estimates), while the zero restrictions provide additional tightening that sharpens the inflationary signal from supply chain shocks.

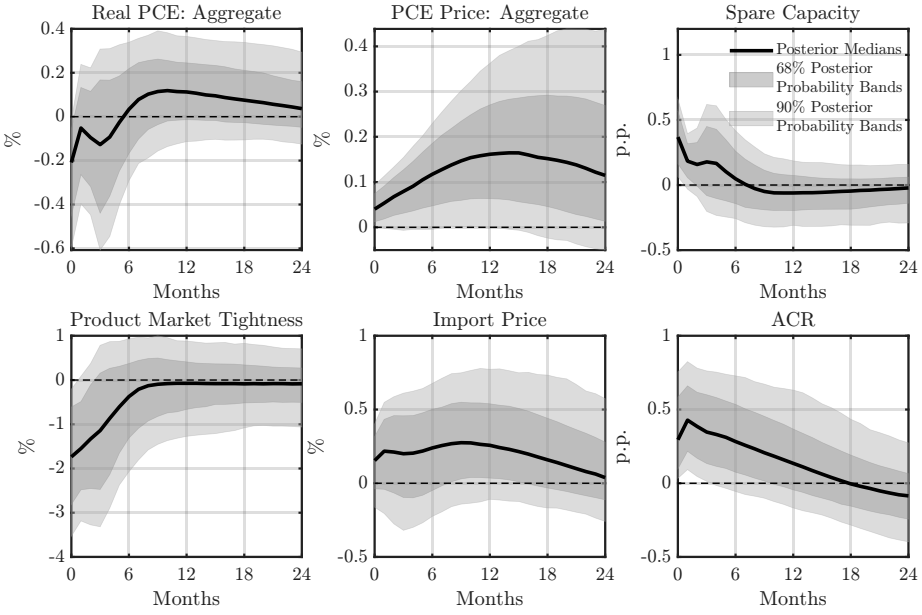


Figure H.4: IRFs to an Adverse Supply Chain Shock: ACR Without Zero Restrictions

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification described in Section III, except that the zero restrictions on the ACR index are omitted.

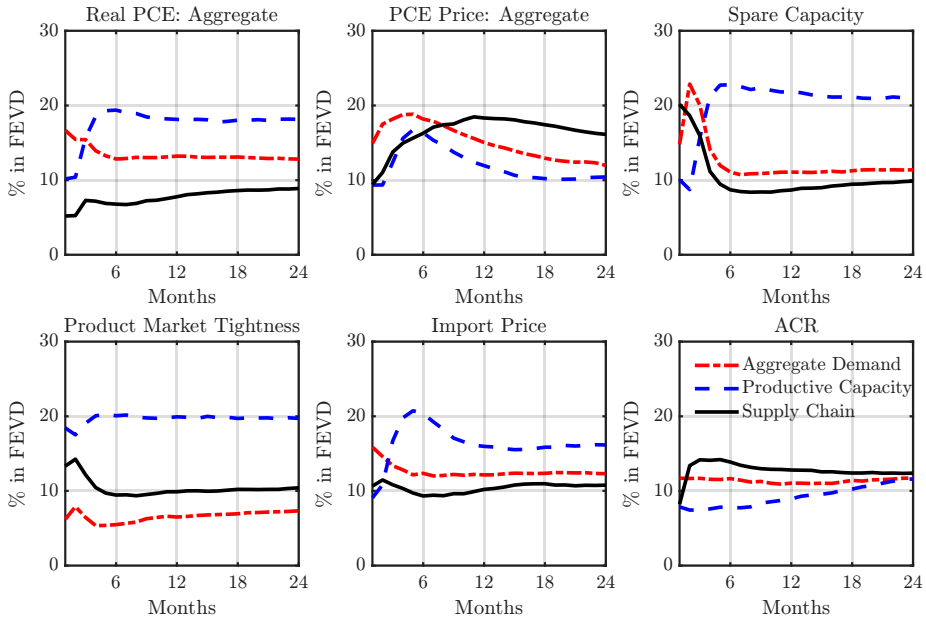


Figure H.5: FEVDs from the SVAR: ACR Without Zero Restrictions

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the zero restrictions on the ACR index are omitted.

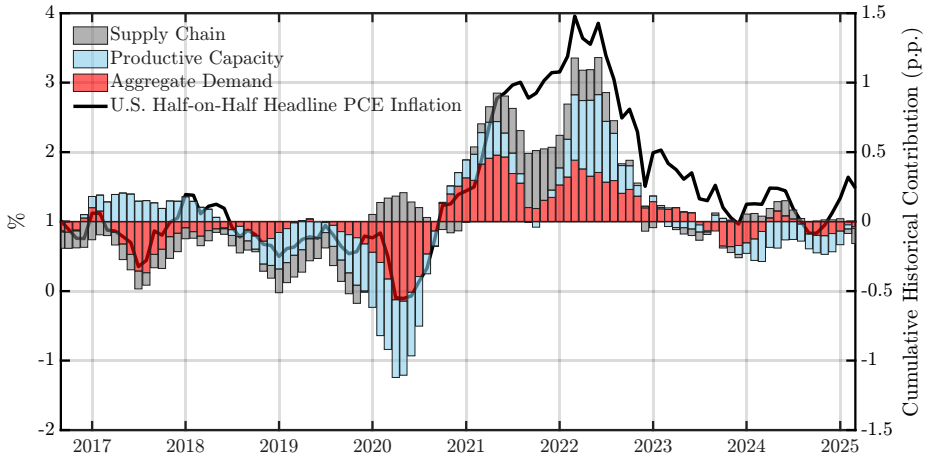


Figure H.6: HD of U.S. HoH Headline PCE Inflation: ACR Without Zero Restrictions

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the zero restrictions on the ACR index are omitted.

H.3. ACT

As introduced in Appendix C.3, we define an alternative port congestion metric, the ACT index. This index measures the average number of hours a container ship waits in a port's anchorage area before docking at a berth, weighted by each port's share of ship visits. Figure H.7 plots the ACT index alongside the ACR index at a monthly frequency. The two series co-move closely and display similar dynamics, except during 2016–2017, when infrastructure upgrades worldwide appear to have focused on the extensive margin (reducing the share of delayed ships) with limited progress on the intensive margin (reducing average delays).[18]

Given the high correlation between ACR and ACT, it is unsurprising that, as shown in Figures H.8–H.10, using the ACT index in the causality assessment yields results that are quantitatively similar to those in the ACR baseline.

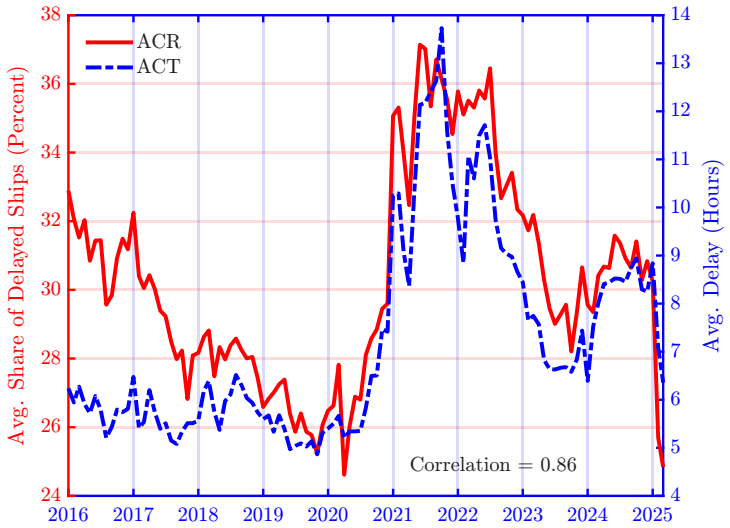


Figure H.7: ACR vs. ACT

Notes. The ACR and the ACT over the sample period January 2016-March 2025 (correlation ≈ 0.86). The ACR is expressed in percentage terms and shown on the left axis, while the ACT is plotted on the right axis in hours. Both indices are seasonally adjusted.

[18] In 2016–2017, capacity expansions (new berths, deeper channels, larger cranes) lowered the likelihood of delays at the extensive margin. By contrast, operational reforms needed to shorten average waiting times at the intensive margin were slower to materialize.

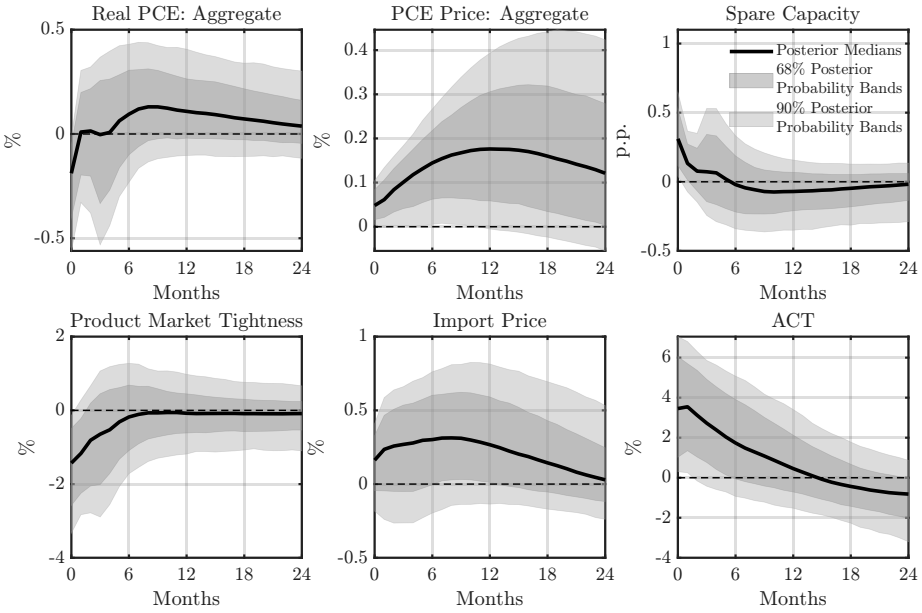


Figure H.8: IRFs to an Adverse Supply Chain Shock: ACT

Notes. The IRFs to a one–standard–deviation adverse supply chain shock are estimated using the baseline SVAR specification in Section III, except that the ACT index replaces the ACR index as the proxy for global supply chain conditions.

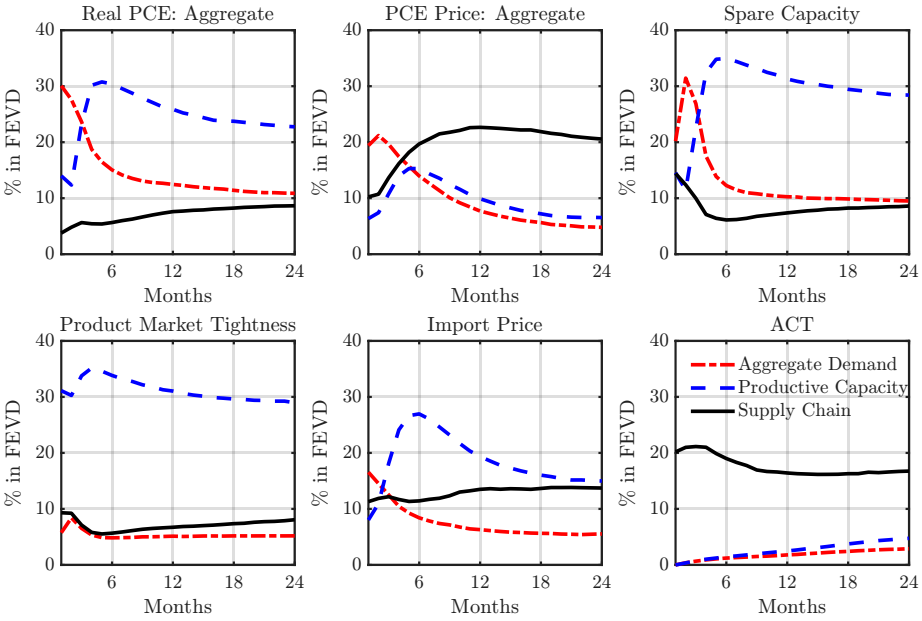


Figure H.9: FEVDs from the SVAR: ACT

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the ACT index replaces the ACR index as the proxy for global supply chain conditions.

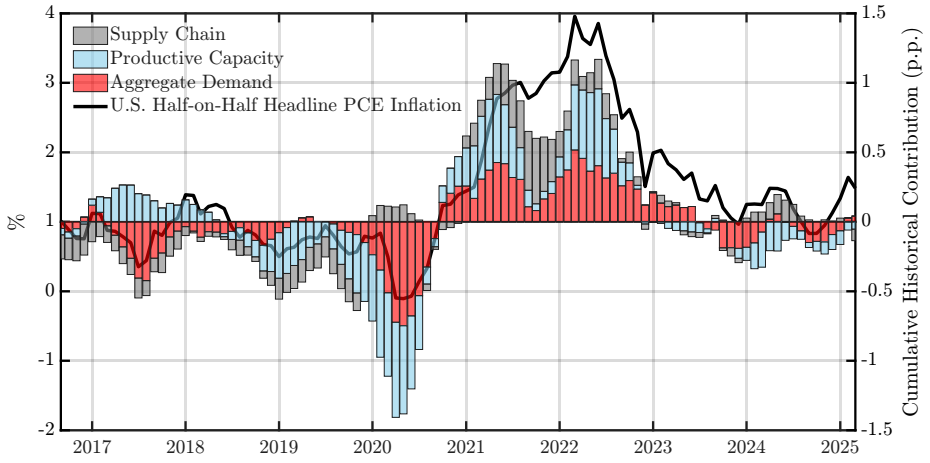


Figure H.10: HD of U.S. HoH Headline PCE Inflation: ACT

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the ACT index replaces the ACR index as the proxy for global supply chain conditions.

H.4. Trans-Pacific ACR

In addition to the global ACR index, we construct a targeted ACR index for the major container ports along the Trans-Pacific route, one of the busiest shipping corridors linking East Asia (primarily China) and the U.S., handling nearly 30 million TEUs annually.[19] We use this index in the SVAR as a robustness check, since congestion along the Trans-Pacific route may exert a disproportionately large influence on the U.S. economy. Shocks identified with this targeted measure may thus reflect supply chain disturbances more specific to the U.S. than those captured by the global ACR index.[20]

Figure H.11 plots the Trans-Pacific ACR index alongside the global ACR index at a monthly frequency. While the two series track each other closely over the sample, the targeted index displays larger short-term swings, particularly during the pandemic.

The estimated effects are similar to those in the ACR baseline. Figure H.12 shows that a supply chain shock generates stagflationary pressures, raises spare capacity, and lowers product

[19]The Trans-Pacific ACR index aggregates congestion at the Port of Busan in South Korea; the Ports of Hanshin and Keihin in Japan; the Ports of Los Angeles and Long Beach in the U.S.; and the Ports of Dalian, Dongguan, Guangzhou Harbor, Hong Kong, Kaohsiung, Ningbo-Zhoushan, Qingdao, Shenzhen, Shanghai, Tianjin, Xiamen, and Yang Shan in China.

[20]Relatedly, Kilian, Nomikos and Zhou (2023) develop a monthly index of container trade volumes to and from North America. Although this measure helps identify shocks to domestic demand and to foreign demand for U.S. manufactured goods, it is not suitable for our purpose of identifying supply chain disturbances.

market tightness. The median inflation responses remain sizable and broadly comparable to the baseline, with the lower bounds of the posterior bands staying above zero within a year of the shock. As shown in Figure H.13, supply chain shocks continue to account for the largest share of the forecast error variance of the PCE price index at medium and long horizons, while capacity shocks remain the dominant driver of the other endogenous variables except for the Trans-Pacific ACR. Finally, the HD in Figure H.14 is nearly indistinguishable from that obtained with the global ACR, indicating that the narratives in Section III regarding the drivers of U.S. headline PCE inflation are robust, even when focusing specifically on congestion at Trans-Pacific ports.

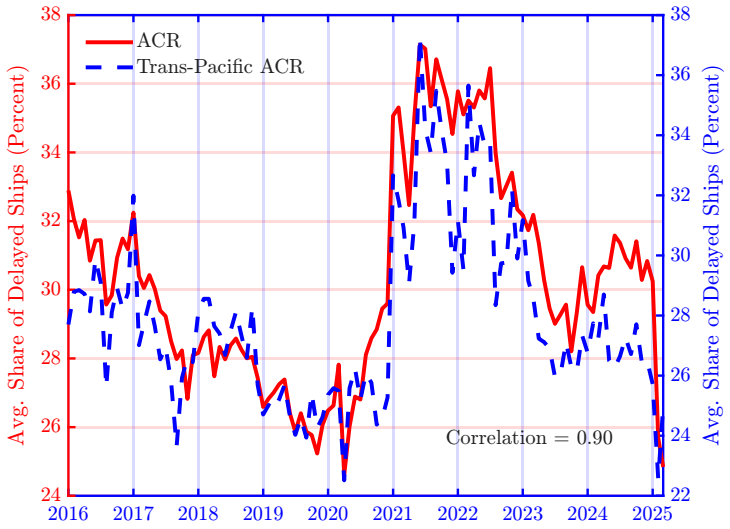


Figure H.11: Global ACR vs. Trans-Pacific ACR

Notes. The global ACR and the targeted ACR for major ports along the Trans-Pacific route over January 2016–March 2025 (correlation ≈ 0.90). Both indices are expressed in percentage terms and seasonally adjusted.

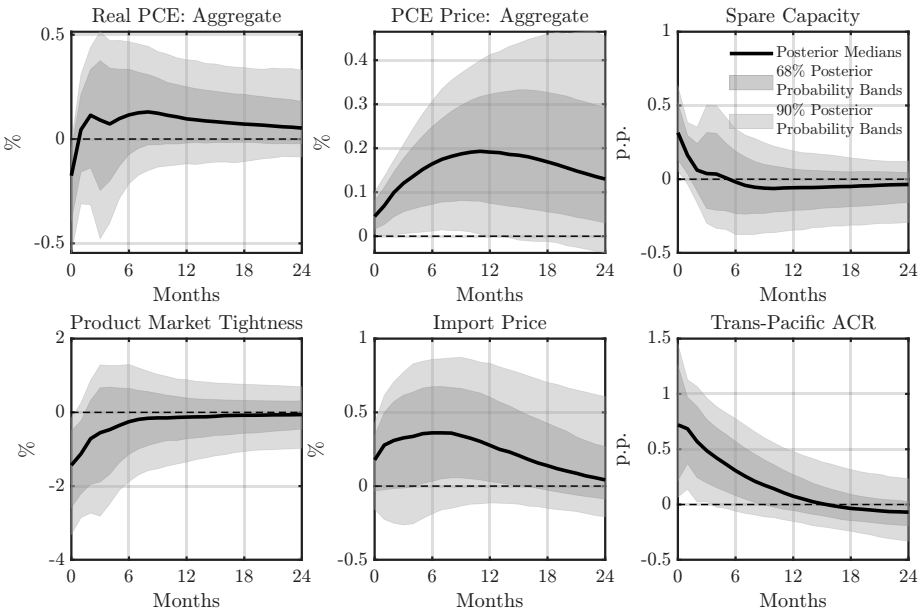


Figure H.12: IRFs to an Adverse Supply Chain Shock: Trans-Pacific ACR

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR described in Section III, except that the targeted ACR index for major Trans-Pacific ports is used as the proxy for the U.S.-specific supply chain conditions.

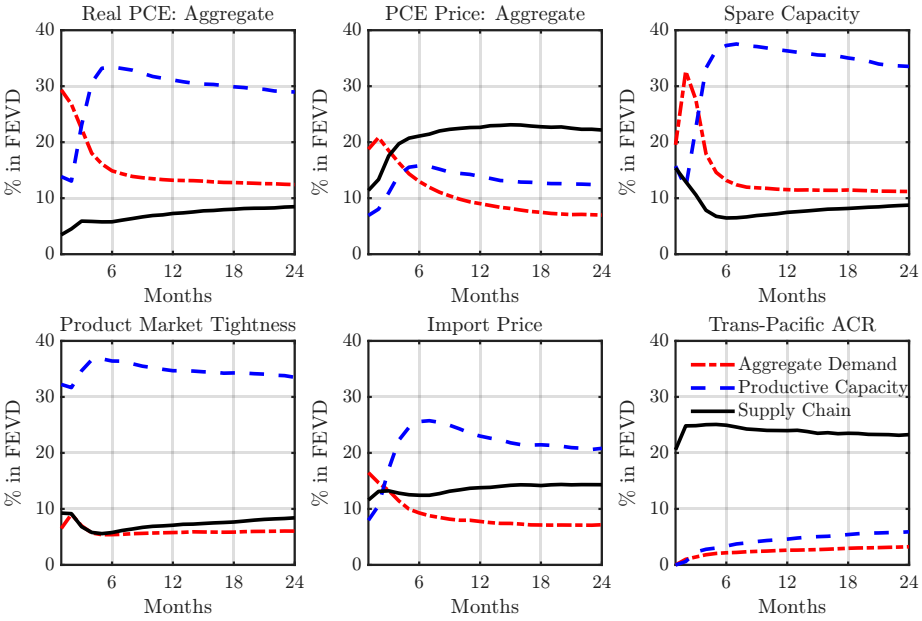


Figure H.13: FEVDs from the SVAR: Trans-Pacific ACR

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable explained by each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR with the same specification as the baseline, except that the targeted ACR index for major Trans-Pacific ports is used as the proxy for the U.S.-specific supply chain conditions.

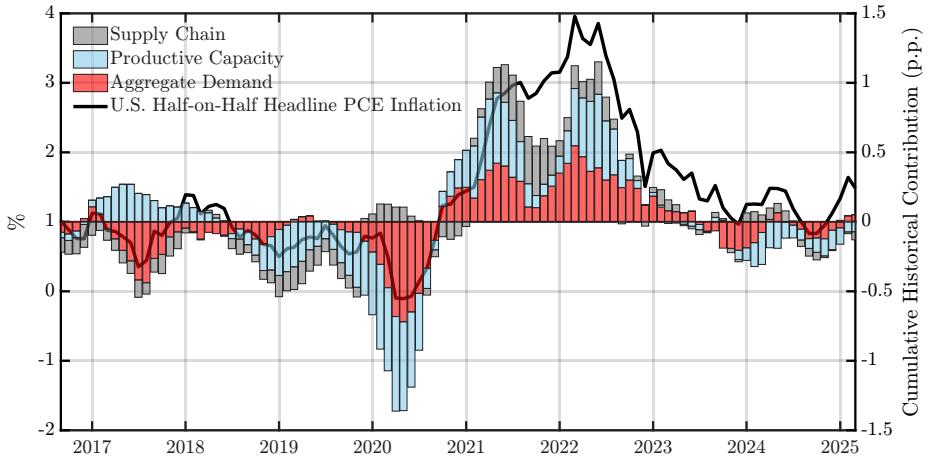


Figure H.14: HD of U.S. HoH Headline PCE Inflation: Trans-Pacific ACR

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the targeted ACR index for major Trans-Pacific ports is used as the proxy for the U.S.-specific supply chain conditions.

H.5. HARPEX

Shipping costs offer a near-real-time read on logistics conditions and are often used to proxy the state of the global supply chain. Within this class, HARPEX (an aggregate of container time-charter rates across vessel sizes) captures the price of shipping capacity rather than congestion itself, and it also enters the New York Fed's GSCPI (Benigno et al., 2022).[21] Because charter rates adjust quickly to shifts in demand, fuel costs, and market expectations, we do not impose the zero restrictions on HARPEX's responses to aggregate demand and productive capacity shocks at horizon $k = 1$. For context, Figure H.15 compares ACR and HARPEX; although the two series co-move, they differ in magnitude and timing, consistent with ACR tracking port congestion and HARPEX reflecting the price of transport services.

Figure H.16 shows that replacing ACR with HARPEX yields a more muted response of the PCE chain-type price index to an adverse supply chain shock: the median effect is small, and the 90% posterior band largely includes zero across horizons. Consistently, the FEVDs in Figure H.17 attribute a larger share of the forecast error variance of PCE prices to aggregate demand at

[21] Finck and Tillmann (2022) also incorporate HARPEX (together with the RWI/ISL container throughput index and the GSCPI) as endogenous variables in SVARs to capture international container shipping and global supply chain conditions. Their identification of a global supply chain shock combines conventional sign restrictions with narrative information (e.g., the 2011 Tōhoku earthquake, the 2021 Suez Canal blockage, and the 2022 Shanghai backlog).

medium horizons, with only a modest portion explained by the supply chain shock, suggesting that HARPEX blends congestion signals with movements driven by demand and other market forces.

The HD in Figure H.18 reinforces this conclusion. During the 2021-22 inflation surge, the contribution from the supply chain shock is present but materially smaller than under the ACR baseline, while demand and capacity account for a greater share. Taken together, the evidence from impulse responses, variance decompositions, and the historical decomposition of U.S. headline PCE inflation indicates that HARPEX primarily reflects broader shipping market dynamics driven by demand and costs, whereas ACR more directly isolates congestion-related pressures.

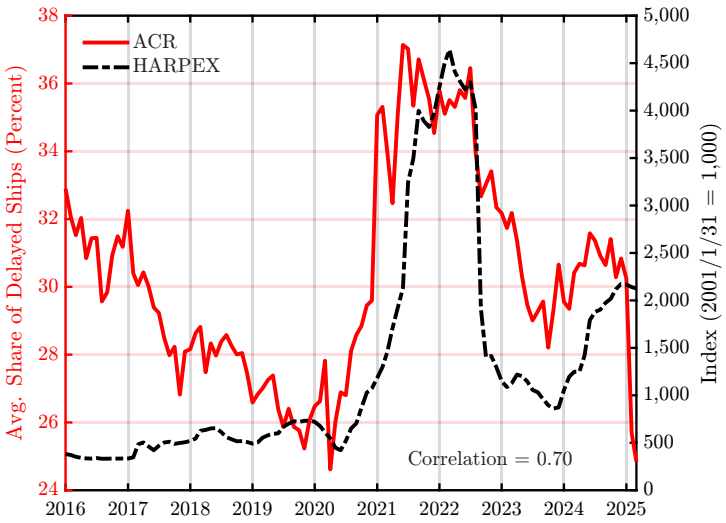


Figure H.15: ACR vs. HARPEX

Notes. The ACR and HARPEX over the sample period January 2016-March 2025 (correlation ≈ 0.70). The ACR is expressed in percentage terms and shown on the left axis, while HARPEX is plotted on the right axis with a baseline value of 1,000 on January 31, 2001. The HARPEX series is published by Harper Peterson and retrieved from the Refinitiv data platform. Both series are seasonally adjusted.

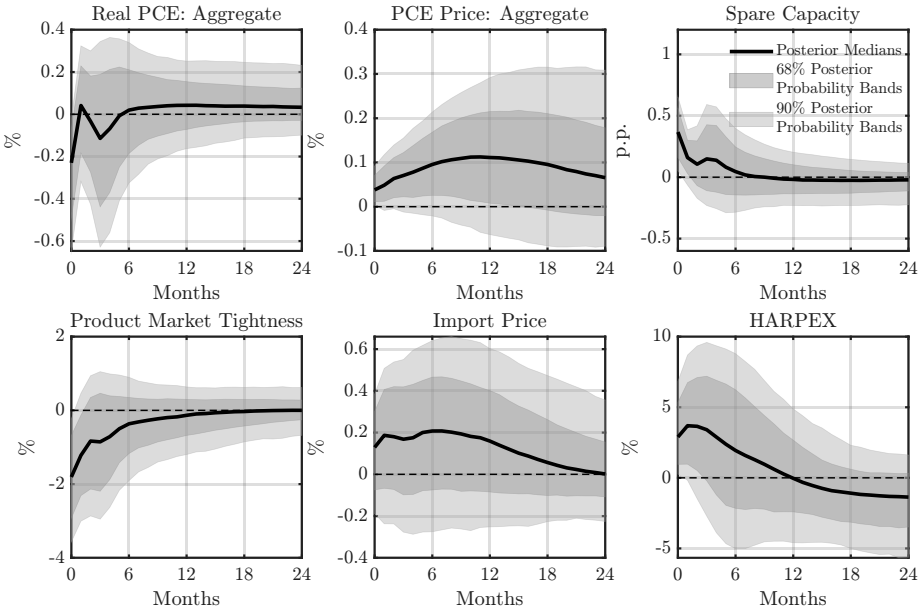


Figure H.16: IRFs to an Adverse Supply Chain Shock: HARPEX

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification described in Section III, except that HARPEX replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

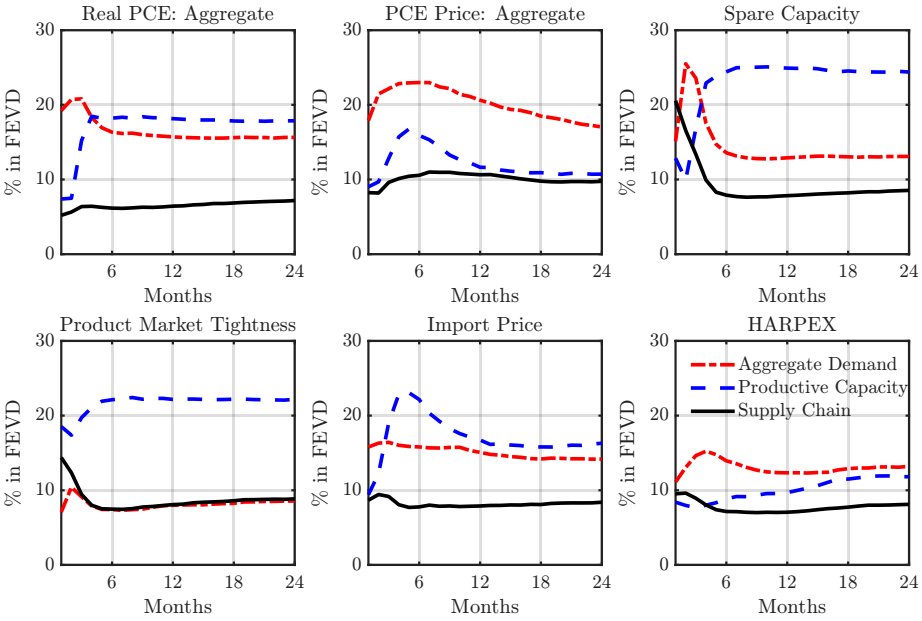


Figure H.17: FEVDs from the SVAR: HARPEX

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that HARPEX replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

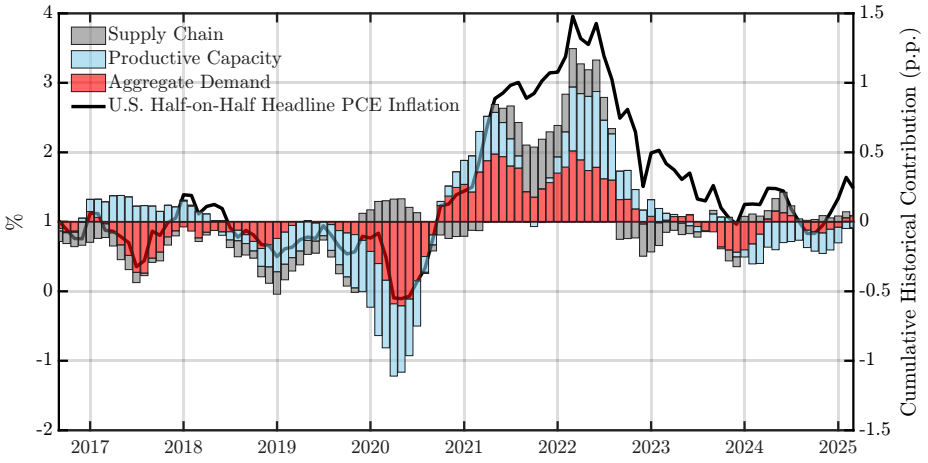


Figure H.18: HD of U.S. HoH Headline PCE Inflation: HARPEX

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that HARPEX replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

H.6. GSCPI

The GSCPI combines cross-border transportation costs and manufacturing PMI subcomponents (e.g., suppliers' delivery times) to infer global supply chain conditions (Benigno et al., 2022). While widely used, these inputs introduce potential measurement errors: transportation costs fluctuate with factors unrelated to supply chain disruptions (such as shifts in fuel prices or shipping demand), PMI indicators reflect subjective managerial assessments that may embed expectations or misperceptions, and longer delivery times can arise from production-side bottlenecks rather than congestion in the logistics network. By contrast, our ACR index is constructed from satellite-based vessel data and directly tracks global port congestion (a first-order manifestation of supply chain disruptions) (Transportation Research Board Executive Committee, 2006).

Figure H.19 compares ACR and the GSCPI over 2016–2025 and shows that, although positively correlated, the two indices diverge in amplitude and timing. During the early pandemic, di Giovanni et al. (2022) interpret the sharp rise in the GSCPI as evidence of China's lockdown and its later decline as reflecting partial reopenings in 2020. The ACR, however, indicates that the lockdown did not generate congestion of a magnitude consistent with severe global disruptions, nor did the reopenings deliver immediate relief at major ports. This contrast foreshadows the

weaker inflation signal obtained when the GSCPI substitutes for ACR.

Figure H.20 shows that the response of the PCE chain-type price index to an adverse supply chain shock is muted: the median effect is small, and the 90% posterior band includes zero throughout, in contrast to the clearer response under ACR with the zero restrictions. Likewise, the FEVDs in Figure H.21 allocate more of the variance in real PCE and in PCE and import prices to aggregate demand at medium horizons, while the supply chain share remains comparatively modest.

The HD in Figure H.22 corroborates this view, showing that the contribution of the supply chain shock to the 2021-22 inflation episode is visible but materially smaller than in the ACR baseline, with demand and capacity accounting for a larger share. Together, these diagnostics point to a diluted imprint of supply chain disturbances when the GSCPI is used in place of ACR.

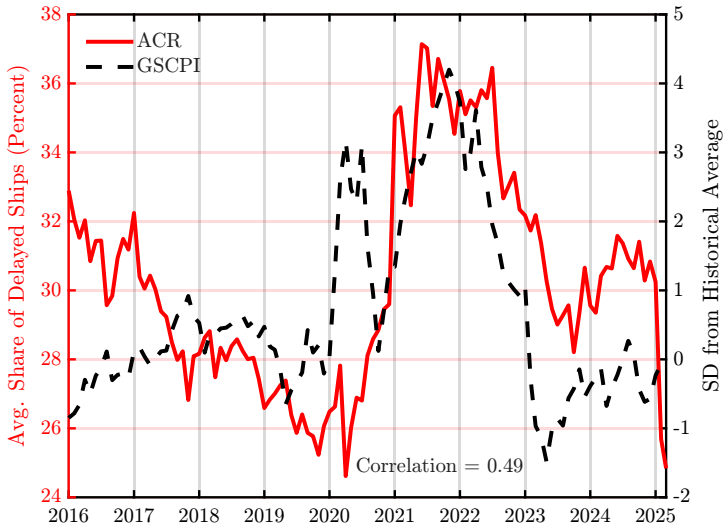


Figure H.19: ACR vs. GSCPI

Notes. The ACR and the New York Fed's GSCPI over January 2016–March 2025 (correlation ≈ 0.49). The ACR is expressed in percentage terms and shown on the left axis, while the GSCPI is plotted on the right axis in standard deviations from its historical average. The GSCPI series is retrieved from the Federal Reserve Bank of New York. Both series are seasonally adjusted.

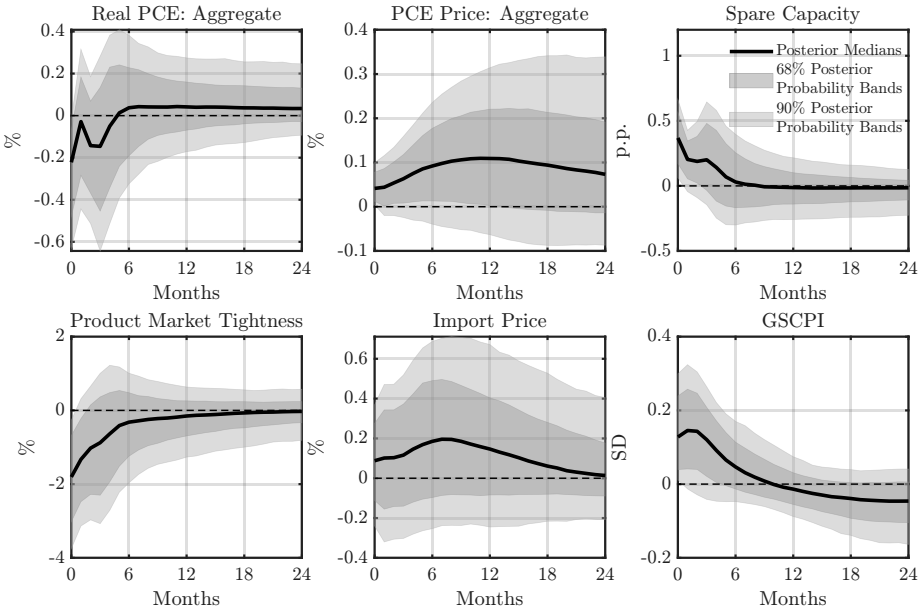


Figure H.20: IRFs to an Adverse Supply Chain Shock: GSCPI

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification described in Section III, except that the GSCPI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

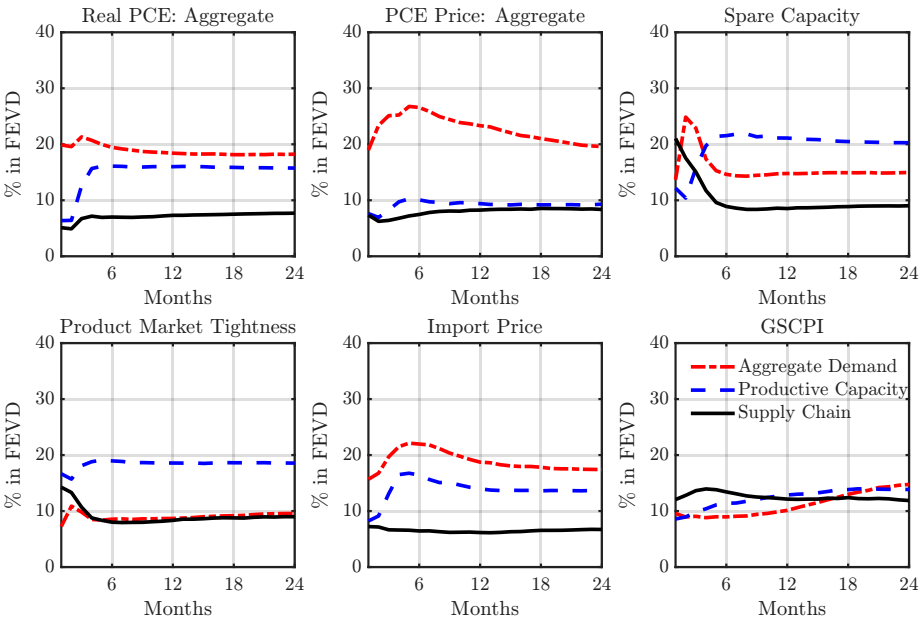


Figure H.21: FEVDs from the SVAR: GSCPI

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the GSCPI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

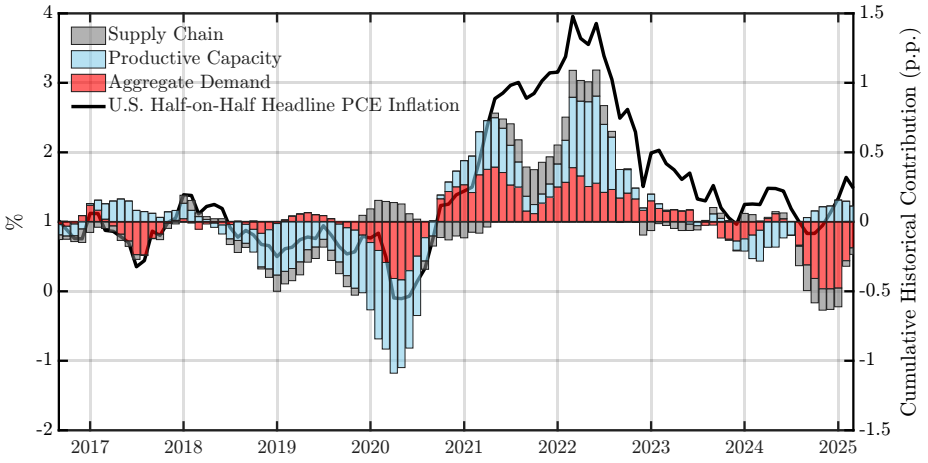


Figure H.22: HD of U.S. HoH Headline PCE Inflation: GSCPI

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the GSCPI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

H.7. Supply Disruptions Index (SDI)

Smirnyagin and Tsyvinski (2022) and Liu, Smirnyagin and Tsyvinski (2024) construct the SDI using the S&P Global Panjiva dataset of U.S. seaborne import records. The index tracks regular consignee-shipper relationships on a quarterly basis and records a disruption when an otherwise active pair becomes inactive for one quarter before resuming. From 2016 to 2025, the SDI has remained relatively stable before the pandemic, rose in early 2020, showed renewed elevation during the 2021-22 inflation episode, and then eased thereafter (Figure H.23).

Turning to the SVAR evidence, Figure H.24 shows that when SDI replaces ACR, the responses of the PCE chain-type price index (and likewise the import price index) to an adverse supply chain shock are less precisely estimated: median effects are small, and the 90% posterior band often includes zero across horizons.

Consistently, the FEVDs in Figure H.25 assign only a small share of fluctuations in PCE and import prices to supply chain shocks at medium horizons. The HD in Figure H.26 likewise shows that supply chain shocks made only a limited contribution to the 2021-22 inflation episode, with demand and capacity playing larger roles. Echoing the results with the GSCPI, substituting SDI for ACR weakens the identification of the inflationary effects of supply chain disturbances.

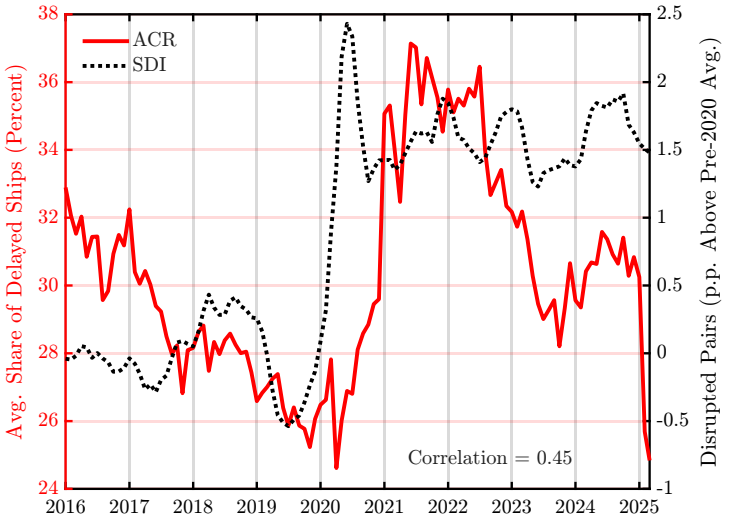


Figure H.23: ACR vs. SDI

Notes. The ACR and the SDI over the sample period January 2016-March 2025 (correlation ≈ 0.45). The ACR is expressed as a percentage and shown on the left axis, while the SDI is plotted on the right axis in percentage points. The SDI series is retrieved from https://www.disruptions.supply (accessed August 3, 2025). Both series are seasonally adjusted.

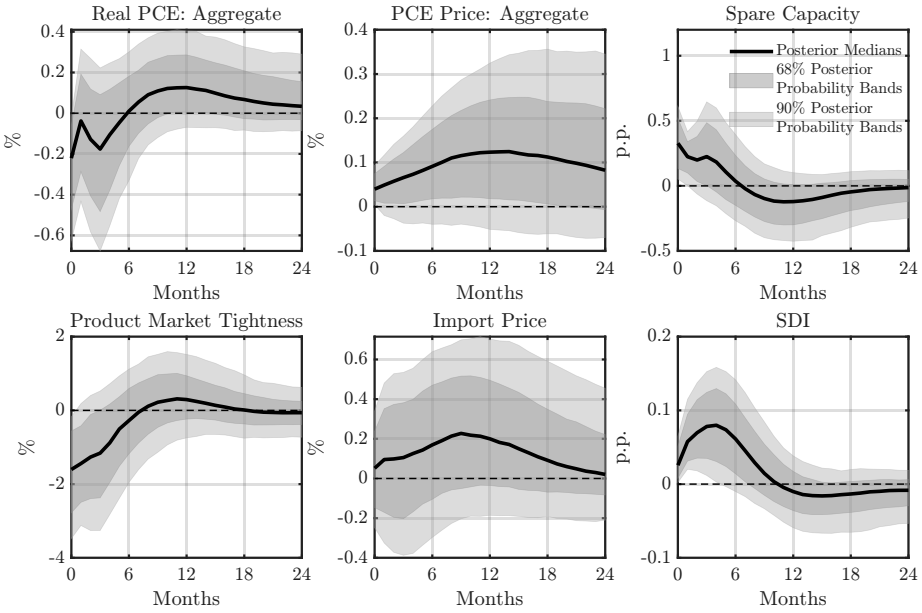


Figure H.24: IRFs to an Adverse Supply Chain Shock: SDI

Notes. The IRFs to a one-standard-deviation adverse supply chain shock are estimated using the baseline SVAR specification in Section III, except that the SDI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

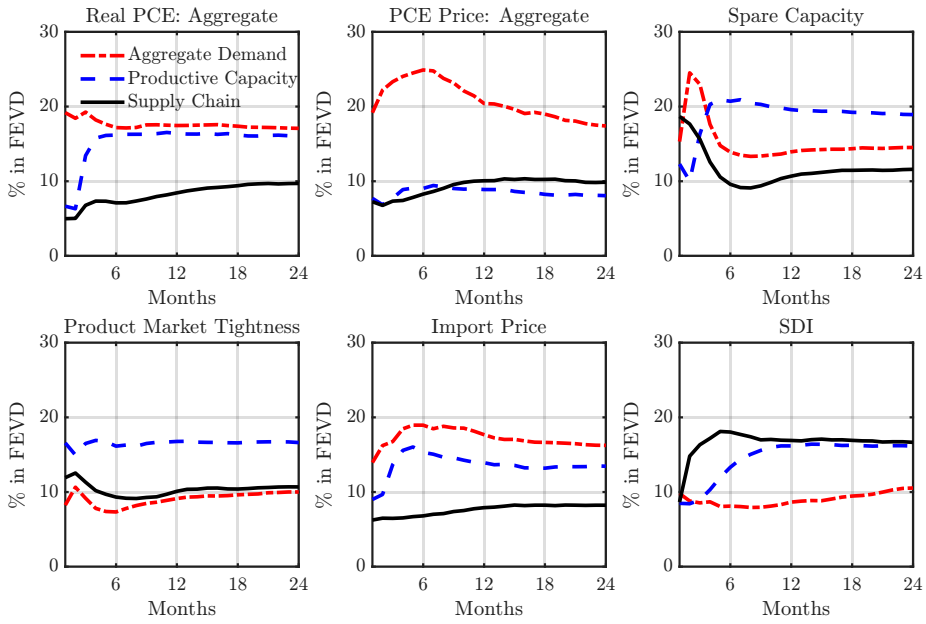


Figure H.25: FEVDs from the SVAR: SDI

Notes. Each line shows the posterior-median share of the forecast error variance of an endogenous variable attributed to each of the three identified structural shocks across horizons. The FEVDs are computed from the Bayesian SVAR identified as in the baseline, except that the SDI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

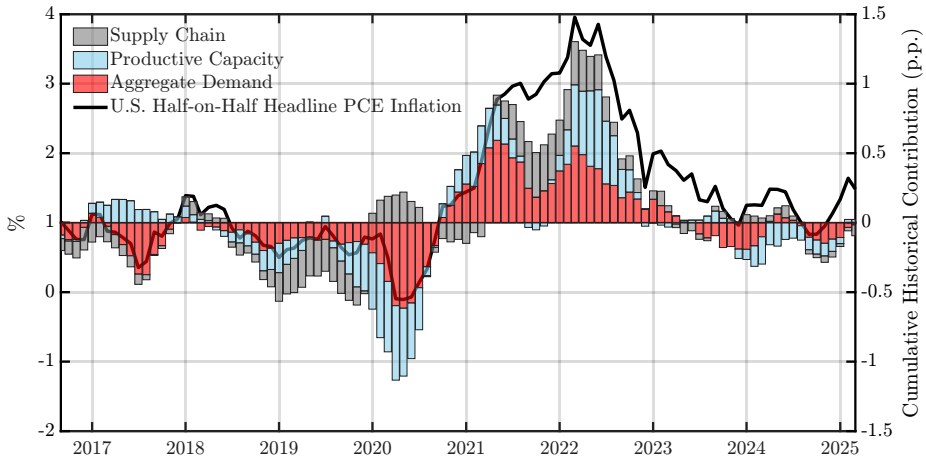


Figure H.26: HD of U.S. HoH Headline PCE Inflation: SDI

Notes. The solid line plots U.S. headline PCE inflation, measured as the HoH growth of the PCE chain-type price index, against the left axis (%). The shaded bars show the corresponding cumulative historical contributions of aggregate demand, productive capacity, and supply chain shocks against the right axis (p.p.). The HD is computed from an identified Bayesian SVAR with the same specification as the baseline, except that the SDI replaces the ACR index as the proxy for global supply chain conditions and its responses to demand and capacity shocks are left unrestricted at horizon $k = 1$.

H.8. Goods Price Responses Across Proxies

We replicate the cross-proxy exercise in Section D using the PCE goods price index instead of the aggregate PCE price index. Figure H.27 reports impulse responses to a one-standard-deviation adverse supply chain shock across eight specifications: no proxy; ACR; ACR with zero restrictions on its responses to demand and capacity shocks at horizon $k = 1$; ACT with the same restrictions; Trans-Pacific ACR with the same restrictions; HARPEX; GSCPI; and SDI. When ACR, ACT, or the Trans-Pacific ACR is combined with these zero restrictions, the median responses are larger and more precisely estimated —posterior bands are tighter, with the 68% interval above zero at nearly all horizons and the 90% interval largely above zero— clearly indicating the inflationary nature of supply chain disturbances and consistent with our theoretical predictions. In contrast, omitting a proxy or using HARPEX, GSCPI, or SDI yields smaller responses with wider bands that generally include zero.

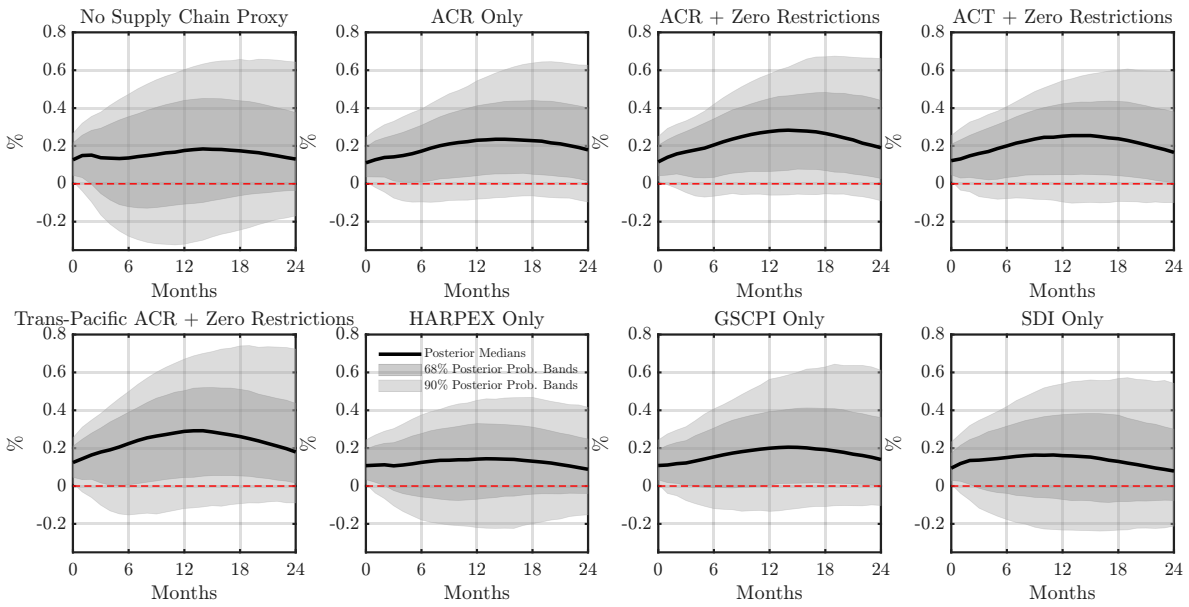


Figure H.27: Goods Price Responses to an Adverse Supply Chain Shock Across Proxies

Notes. Posterior-median impulse responses of the U.S. PCE goods price index, with 68% and 90% posterior probability bands, to a one-standard-deviation adverse supply chain shock across eight SVAR specifications: (i) no supply chain proxy; (ii) ACR; (iii) ACR with zero restrictions at horizon $k = 1$ on its responses to aggregate demand and productive capacity shocks; (iv) ACT with the same zero restrictions; (v) Trans-Pacific ACR with the same zero restrictions; (vi) HARPEX; (vii) GSCPI; and (viii) SDI. Aside from the choice of proxy, all eight specifications follow the baseline model in Section III, except that real PCE and the aggregate PCE price index are replaced by goods PCE and the PCE goods price index. The sign restrictions in Restrictions 1-3 are imposed to identify the adverse supply chain shock in all specifications except (i), where the positive restriction on the ACR response is omitted.

The FEVD results in Figure H.28 show that the congestion indices, together with the zero restrictions motivated by our domain knowledge, assign a larger share of the forecast error variance of PCE goods prices to supply chain disturbances at medium horizons. In contrast, with alternative proxies, the attribution shifts toward aggregate demand. These patterns mirror the aggregate-price results and highlight the identification gains from combining congestion-based proxies with zero restrictions.

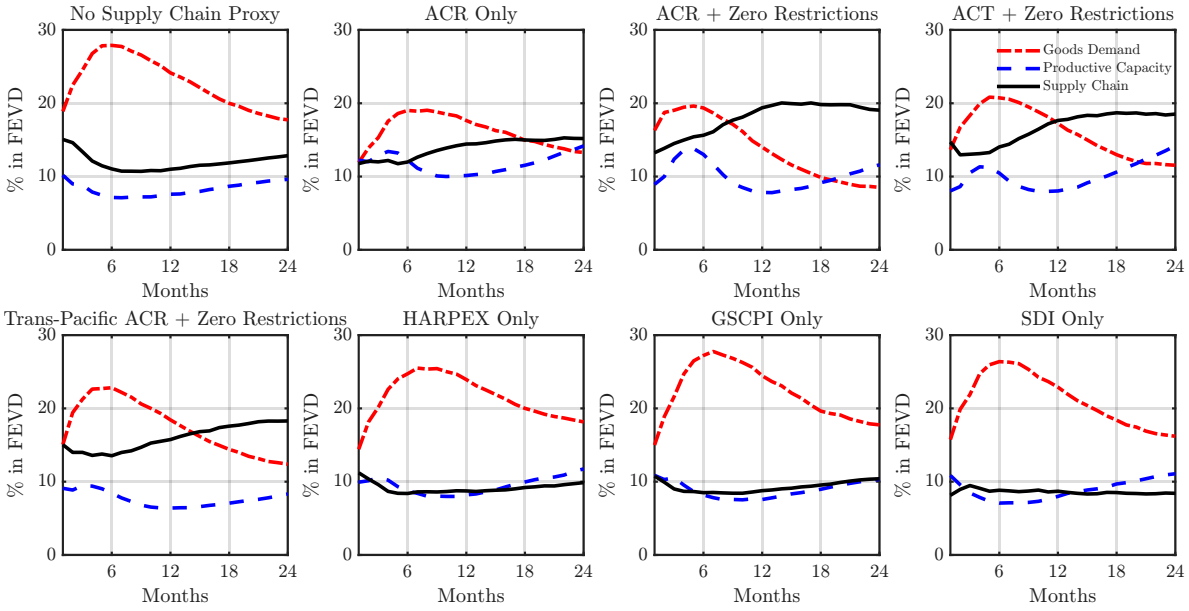


Figure H.28: FEVD of the U.S. PCE Goods Price Index Across Proxies

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

Appendix References

Allen, Treb. 2014. "Information Frictions in Trade." *Econometrica*, 82: 2041–2083.

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

Bai, Xiwen, Zhongjun Ma, Yao Hou, Yiliang Li, and Dong Yang. 2023. "A Data-Driven Iterative Multi-Attribute Clustering Algorithm and Its Application in Port Congestion Estimation." *IEEE Transactions on Intelligent Transportation Systems*, 24: 12026–12037.

Baqaee, David, and Emmanuel Farhi. 2019. "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem." *Econometrica*, 87(4): 1155–1203.

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

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

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

Benguria, Felipe. 2021. "The Matching and Sorting of Exporting and Importing Firms: Theory and Evidence." *Journal of International Economics*, 131: 103430.

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

Bils, Mark, Yongsung Chang, and Sun-Bin Kim. 2011. "Worker Heterogeneity and Endogenous Separations in a Matching Model of Unemployment Fluctuations." *American Economic Journal: Macroeconomics*, 3: 128–154.

Birant, Derya, and Alp Kut. 2007. "ST-DBSCAN: An Algorithm for Clustering Spatial–Temporal Data." *Data & Knowledge Engineering*, 60: 208–221.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. "Geography, Transportation, and Endogenous Trade Costs." *Econometrica*, 88(2): 657–691.

Chaney, Thomas. 2014. "The Network Structure of International Trade." *American Economic Review*, 104(11): 3600–3634.

Chow, Gregory C, and An-loh Lin. 1971. "Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series." *Review of Economics and Statistics*, 53: 372–375.

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

Du, Yuquan, Qiushuang Chen, Jasmine Siu Lee Lam, Ya Xu, and Jin Xin Cao. 2015. "Modeling the Impacts of Tides and the Virtual Arrival Policy in Berth Allocation." *Transportation Science*, 49(4): 939–956.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, 70: 1741–1779.

Ester, Martin, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. 1996. "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise." 226–231. AAAI Press.

Fernández-Villaverde, Jesú, Federico Mandelman, Yang Yu, and Francesco Zanetti. 2024. "Search Complementarities, Aggregate Fluctuations, and Fiscal Policy." *Review of Economic Studies*, rdae053.

Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson. 2007. "ABCs (and Ds) of Understanding VARs." *American Economic Review*, 97(3): 1021–1026.

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

Fujita, Shigeru, and Garey Ramey. 2012. "Exogenous Versus Endogenous Separation." *American Economic Journal: Macroeconomics*, 4: 68–93.

Giacomini, Raffaella, and Toru Kitagawa. 2021. "Robust Bayesian Inference for Set-Identified Models." *Econometrica*, 89: 1519–1556.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, 95(1): 161–182.

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

Kilian, Lutz, Nikos Nomikos, and Xiaoqing Zhou. 2023. "Container Trade and the U.S. Recovery." *International Journal of Central Banking*, 19(1): 417–450.

Krolikowski, Pawel M, and Andrew H McCallum. 2021. "Goods-Market Frictions and International Trade." *Journal of International Economics*, 129: 103411.

Lenoir, Clémence, Julien Martin, and Isabelle Mejean. 2022. "Search Frictions in International Goods Markets." *Journal of the European Economic Association*, 21(1): 326–366.

Li, Chen, Xiangtong Qi, and Dongping Song. 2016. "Real-Time Schedule Recovery in Liner Shipping Service With Regular Uncertainties and Disruption Events." *Transportation Research Part B: Methodological*, 93: 762–788.

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

Li, Yiliang, Shengyao Xu, Senyi Xia, Zhongjun Ma, and Xiwen Bai. 2025. "Effective Global Maritime Supply: Determinants, Dynamics, and Economic Impacts." Tsinghua University Working Paper.

Melitz, Jacques, and Farid Toubal. 2014. "Native Language, Spoken Language, Translation and Trade." *Journal of International Economics*, 93: 351–363.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica*, 71: 1695–1725.

Menzio, Guido, and Shouyong Shi. 2011. "Efficient Search on the Job and the Business Cycle." *Journal of Political Economy*, 119: 468–510.

Naudé, Wim, and Marianne Matthee. 2011. "The Impact of Transport Costs on New Venture Internationalisation." *Journal of International Entrepreneurship*, 9: 62–89.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

Notteboom, Theo E. 2006. "The Time Factor in Liner Shipping Services." *Maritime Economics & Logistics*, 8: 19–39.

Rodrigue, Jean-Paul. 2020. *The Geography of Transport Systems.* . 5 ed., Routledge.

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

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

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

Wong, Woan Foong. 2022. "The Round Trip Effect: Endogenous Transport Costs and International Trade." *American Economic Journal: Applied Economics*, 14: 127–166.

Data References

Bauer, Michael D., and Eric T. Swanson. 2025. "Monetary Policy Surprises." *Federal Reserve Bank of San Francisco*, https://www.frbsf.org/research-and-insights/data-and-indicators/monetary-policy-surprises/ (accessed July 5, 2025).

Board of Governors of the Federal Reserve System. 2025*a*. "Federal Funds Effective Rate [FEDFUNDS]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/FEDFUNDS (accessed July 10, 2025).

Board of Governors of the Federal Reserve System. 2025*b*. "Industrial Production: Total Index [INDPRO]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/INDPRO (accessed July 10, 2025).

Clarksons Research. 2025*a*. "Containership Average Speed - Knots [8852030]." *Shipping Intelligence Network*, accessed July 6, 2025.

Clarksons Research. 2025*b*. "Total Containerships - % Idle, Number of Ships [545702]." *Shipping Intelligence Network*, accessed July 6, 2025.

European Commission. 2025. "Germany Capacity Utilization." *Trading Economics*, https://tradingeconomics.com/germany/capacity-utilization (accessed July 8, 2025).

Federal Reserve Bank of New York. 2025. "Global Supply Chain Pressure Index [GSCPI]." https://www.newyorkfed.org/research/policy/gscpi#/overview (accessed July 9, 2025).

Federal Reserve Bank of St. Louis. 2025. "Spot Crude Oil Price: West Texas Intermediate [WTISPLC]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/WTISPLC (accessed July 13, 2025).

Federal Statistical Office. 2025. "Germany Industrial Production MoM." *Trading Economics*, https://tradingeconomics.com/germany/industrial-production-mom (accessed July 8, 2025).

Harper Petersen & Co. 2025. "HARPEX Charter Rate Index [CHT-IDX-HARPX]." *LSEG Data & Analytics*, accessed July 5, 2025.

International Monetary Fund. 2025. "Global Price of Brent Crude [POILBREUSDM]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/POILBREUSDM (accessed July 7, 2025).

Liu, Ernest, Vladimir Smirnyagin, and Aleh Tsyvinski. 2025. "US Supply Disruptions Index [US SDI]." https://www.disruptions.supply/home (accessed July 9, 2025).

Ministry of Economy, Trade and Industry. 2025*a*. "Japan Capacity Utilization." *Trading Economics*, https://tradingeconomics.com/japan/capacity-utilization (accessed July 8, 2025).

Ministry of Economy, Trade and Industry. 2025*b*. "Japan Industrial Production MoM." *Trading Economics*, https://tradingeconomics.com/japan/industrial-production-mom (accessed July 8, 2025).

National Bureau of Statistics of China. 2025*a*. "China Industrial Capacity Utilization." *Trading Economics*, https://tradingeconomics.com/china/capacity-utilization (accessed July 8, 2025).

National Bureau of Statistics of China. 2025*b*. "China Industrial Production MoM." *Trading Economics*, https://tradingeconomics.com/china/industrial-production-mom (accessed July 8, 2025).

National Institute of Statistics and Geography. 2025*a*. "Mexico Capacity Utilization." *Trading Economics*, https://tradingeconomics.com/mexico/capacity-utilization (accessed July 8, 2025).

National Institute of Statistics and Geography. 2025*b*. "Mexico Industrial Production MoM." *Trading Economics*, https://tradingeconomics.com/mexico/industrial-production-mom (accessed July 8, 2025).

Statistics Canada. 2025*a*. "Canada Capacity Utilization." *Trading Economics*, https://tradingeconomics.com/canada/capacity-utilization (accessed July 8, 2025).

Statistics Canada. 2025*b*. "Canada Industrial Production MoM." *Trading Economics*, https://tradingeconomics.com/canada/industrial-production-mom (accessed July 8, 2025).

U.S. Bureau of Economic Analysis. 2025*a*. "Gross Domestic Product: Implicit Price Deflator [GDPDEF]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/GDPDEF (accessed July 10, 2025).

U.S. Bureau of Economic Analysis. 2025*b*. "Personal Consumption Expenditures: Chain-type Price Index [PCEPI]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/PCEPI (accessed July 12, 2025).

U.S. Bureau of Economic Analysis. 2025*c*. "Personal consumption expenditures: Goods (chain-type price index) [DGDSRG3M086SBEA]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/DGDSRG3M086SBEA (accessed July 12, 2025).

U.S. Bureau of Economic Analysis. 2025*d*. "Personal consumption expenditures: Services (chain-type price index) [DSERRG3M086SBEA]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/DSERRG3M086SBEA (accessed July 12, 2025).

U.S. Bureau of Economic Analysis. 2025*e*. "Real Gross Domestic Product [GDPC1]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/GDPC1 (accessed July 10, 2025).

U.S. Bureau of Economic Analysis. 2025*f*. "Real Personal Consumption Expenditures: Goods [DGDSRX1]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/DGDSRX1 (accessed July 12, 2025).

U.S. Bureau of Economic Analysis. 2025*g*. "Real Personal Consumption Expenditures [PCEC96]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/PCEC96 (accessed July 12, 2025).

U.S. Bureau of Economic Analysis. 2025*h*. "Real Personal Consumption Expenditures: Services [PCESC96]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/PCESC96 (accessed July 12, 2025).

U.S. Bureau of Labor Statistics. 2025*a*. "Average Hourly Earnings of All Employees, Total Private [CES0500000003]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/CES0500000003 (accessed July 13, 2025).

U.S. Bureau of Labor Statistics. 2025*b*. "Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/CPIAUCSL (accessed July 10, 2025).

U.S. Bureau of Labor Statistics. 2025*c*. "Import Price Index (End Use): All Commodities [IR]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/IR (accessed July 10, 2025).

U.S. Bureau of Labor Statistics. 2025*d*. "Producer Price Index by Commodity: Final Demand: Finished Goods [WPSFD49207]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/WPSFD49207 (accessed July 10, 2025).

U.S. Bureau of Labor Statistics. 2025*e*. "Producer Price Index by Industry: Total Manufacturing Industries [PCUOMFGOMFG]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/PCUOMFGOMFG (accessed July 10, 2025).

U.S. Census Bureau. 2025*a*. "Manufacturers' New Orders: Total Manufacturing [AMTMNO]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/AMTMNO (accessed July 17, 2025).

U.S. Census Bureau. 2025*b*. "U.S. Imports of Goods by Customs Basis from Canada [IMPCA]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/IMPCA (accessed July 8, 2025).

U.S. Census Bureau. 2025*c*. "U.S. Imports of Goods by Customs Basis from China [IMPCH]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/IMPCH (accessed July 8, 2025).

U.S. Census Bureau. 2025*d*. "U.S. Imports of Goods by Customs Basis from Germany [IMPGE]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/IMPGE (accessed July 8, 2025).

U.S. Census Bureau. 2025*e*. "U.S. Imports of Goods by Customs Basis from Japan [IMPJP]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/IMPJP (accessed July 8, 2025).

U.S. Census Bureau. 2025*f*. "U.S. Imports of Goods by Customs Basis from Mexico [IMPMX]." *FRED, Federal Reserve Bank of St. Louis*, https://fred.stlouisfed.org/series/IMPMX (accessed July 8, 2025).

World Bank. 2025*a*. "Canada - Industrial Production, Constant US$, Seas. Adj." *Trading Economics*, https://tradingeconomics.com/canada/industrial-production-constant-us$-seas-adj--wb-data.html (accessed July 7, 2025).

World Bank. 2025*b*. "China - Industrial Production, Constant US$." *Trading Economics*, https://tradingeconomics.com/china/industrial-production-constant-us$-wb-data.html (accessed July 7, 2025).

World Bank. 2025*c*. "Germany - Industrial Production, Constant US$, Seas. Adj." *Trading Economics*, https://tradingeconomics.com/germany/industrial-production-constant-us$-seas-adj--wb-data.html (accessed July 7, 2025).

World Bank. 2025*d*. "Japan - Industrial Production, Constant US$, Seas. Adj." *Trading Economics*, https://tradingeconomics.com/japan/industrial-production-constant-us$-seas-adj--wb-data.html (accessed July 7, 2025).

World Bank. 2025*e*. "Mexico - Industrial Production, Constant US$, Seas. Adj." *Trading Economics*, https://tradingeconomics.com/mexico/industrial-production-constant-us$-seas-adj--wb-data.html (accessed July 7, 2025).