NBER WORKING PAPER SERIES

SHIPPING TO AMERICA

Xiwen Bai
Jesús Fernández-Villaverde
Yiliang Li
Ricardo Marto
Francesco Zanetti

Working Paper 35808
http://www.nber.org/papers/w35808

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2026

Nothing to disclose . The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li, Ricardo Marto, and Francesco Zanetti. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Shipping to America
Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li, Ricardo Marto, and Francesco Zanetti
NBER Working Paper No. 35808
September 2026
JEL No. E23, F12, L91, R40

ABSTRACT

We study the macroeconomic and trade-policy implications of disruptions to U.S.-bound shipping routes. Standard models treat them as iceberg-cost shocks, conflating the shock with the response to it. Using satellite vessel-tracking data, we construct route-level measures of potential and effective capacity for all U.S.-bound container ships from 2016 to 2025. Utilization losses in recent disruptions ran 20 to 40 percentage points, and began months before port congestion became visible. We embed these measures in a general equilibrium model in which firms reallocate a common fleet without internalizing the congestion they create and price above marginal cost, while importers' sourcing responds to route profitability. The reallocation triggered by a disruption then has first-order welfare effects, and the route's Domar weight is not a sufficient statistic for its welfare cost. The 2021 West Coast crisis and the 2023-2024 Red Sea attacks cost 0.69% and 0.35% of output. Naval protection of Red Sea shipping generated benefits of 0.04-0.08% of output at a fiscal cost of 0.02%. Tariffs decongest the routes they tax, offsetting or even reversing their conventional welfare cost.

Xiwen Bai
Tsinghua University
xiwenbai@mail.tsinghua.edu.cn

Ricardo Marto
Federal Reserve Bank of St. Louis
rmarto@sas.upenn.edu

Jesús Fernández-Villaverde
University of Pennsylvania
Department of Economics
and CEPR
and also NBER
jesusfv@econ.upenn.edu

Francesco Zanetti
University of Oxford
Department of Economics
francesco.zanetti@economics.ox.ac.uk

Yiliang Li
Peking University
HSBC Business School
yiliang_li@outlook.com

1. Introduction

What are the macroeconomic consequences and trade-policy implications of disruptions to U.S.-bound shipping routes? The 2021–2022 COVID-related supply chain crisis, the 2021 Suez Canal blockage, the 2023–2024 Red Sea crisis, and the 2025–2026 tensions in the Strait of Hormuz are some recent examples of a phenomenon that is becoming increasingly frequent.

Standard models treat a shipping disruption as a shock to the iceberg cost of trade: a larger fraction of every shipment melts in transit. But ships are not icebergs. A ship waiting at anchor burns almost no fuel and is back in service as soon as the port clears. By contrast, a ship rerouted around a blocked chokepoint burns far more fuel and the fleet is short of available tonnage for many weeks. For instance, when Houthi attacks closed the Red Sea route, we had a shock. But the fleet's rerouting around the Cape of Good Hope and the shift of importers toward the routes still open were equilibrium responses that carried the loss to corridors the Houthis never touched. An iceberg cost adds the shock and the endogenous response together and calls the sum exogenous.

Adding the two together also explains why standard models find that shipping disruptions, and trade disruptions in general, barely matter. In these models, the response is free: when one route becomes expensive, capacity and sourcing reallocate, and the envelope theorem makes that reallocation costless at first order. Since the disruption matters only in proportion to the affected route's Domar weight (Hulten, 1978), we get small numbers almost by construction. Hulten's logic does not survive when we consider the fleet's endogenous response. In the short run, the fleet is a common pool that operates in a distorted shipping network: firms do not internalize the congestion they create and, thanks to their market power, price above marginal cost, while importers earn sourcing rents. The corresponding reallocation moves capacity between uses with different social returns and carries first-order welfare effects of its own.

To separate shocks from endogenous responses, we construct route-level measures of *potential* and *effective* shipping capacity from high-frequency Automatic Identification System (AIS) records matched to vessel-level characteristics from the Lloyd's List registry, covering the universe of U.S.-bound container ships from January 2016 through March 2025. Potential capacity measures the services the deployed fleet could provide under benchmark operating conditions; effective capacity subtracts the losses from port congestion, Suez and Panama Canal bottlenecks, rerouting, and slow steaming. Their ratio is a measure of shipping utilization comparable across routes and over time.

Expanding potential capacity does not always increase effective shipping services. U.S.-bound potential capacity rose more than 60% above its pre-pandemic level by early 2022 as firms redirected vessels from Africa, Latin America, and the Caribbean toward the U.S. But aggregate shipping utilization fell from about 89% before the pandemic to below 70%: at the

trough, nearly a third of the deployed capacity delivered no services at all. The decline was concentrated on the Asia–West Coast route, where port waiting and slower operations absorbed nearly half of potential capacity. The fall had already begun by November 2020, months before congestion at Los Angeles and Long Beach became visible in mid-2021. The same thing happened in the Red Sea: measured utilization fell again sharply because rerouting via the Cape of Good Hope added 600–800 nautical miles to voyages from the Middle East and 1,000–1,250 to voyages from Asia.

These observations motivate a model in which vessel deployment and the conversion of deployed capacity into services are endogenous. Shipping firms allocate a common fleet across routes and choose sailing speed, loading intensity, and prices, while importers' sourcing decisions respond to route profitability. Congestion, markups, and sourcing rents are distortions in the sense of Baqaee and Farhi (2019, 2020), and they break Hulten (1978)'s theorem: the reallocation of capacity and sourcing triggered by a disruption has first-order welfare effects throughout the network, so neither the Domar weight nor the domestic expenditure share of Arkolakis et al. (2012) is a sufficient statistic.

Departures from the Hulten benchmark are large and may change sign. We discipline the model using route-level utilization, deployed capacity, speed, and import shares, and simulate it monthly from 2016 to 2025. In our calibrated economy, the output elasticity to an Asia–West Coast disruption is 1.3 times the route's Domar weight, because the disruption also drives importers out of the market and reinforces the direct loss of imports. A disruption on the Middle East or South America routes instead marginally raises output as it shifts capacity toward the more valuable Asian routes.

The 2021 West Coast port crisis cost 0.69% of output. That number comes from a counterfactual holding Asia–West Coast shipping utilization at its January 2020 value, roughly 1.5 times its end-2021 trough: effective capacity rises by 60% and the import share of GDP by 1.1 percentage points (p.p.) at the trough; the retail price of Asian goods falls by 23% and the aggregate price index by 1.1%. We measure the welfare cost as an equivalent variation. The same counterfactual implies that the disruptions added 1.1 p.p. to CPI inflation by December 2021. An analogous counterfactual for the 2023–2024 Houthi attacks, holding utilization on the affected routes at August 2023 values, yields a cost of 0.35% of output.

Three policy applications exploit the same mechanism. First, we evaluate scenarios in which Red Sea utilization would have kept deteriorating absent naval protection. The Navy's protection generated a welfare gain of 0.06% of output in our central scenario, against a fiscal cost of 0.008–0.03% of GDP. Second, tariffs cut not only import demand but also congestion on the routes serving the taxed region. This decongestion force attenuates the conventional price and welfare effects of a tariff. In our March 2025 calibration, the force is so strong that it reverses

them. Third, we quantify the gains from port infrastructure investments to reduce congestion. Our results imply that the macroeconomic cost of a shipping disruption cannot be inferred from the size of the trade flow it affects. What matters instead is how congested the network already is, something an iceberg cost misses. Our route-level measures let us see that congestion directly, and congestion can make the effects of trade policy look quite different from what one might expect.

This paper connects three strands of research. First, it contributes to the literature on international transportation and the industrial organization of shipping. Classic accounts emphasize the technology and organization of maritime transportation (Stopford, 2008), while recent work studies how fleet investment, search frictions, ballast trips, and congestion shape shipping capacity and freight rates (Kalouptsidi, 2014, Greenwood and Hanson, 2015, Dunn and Leibovici, 2023, Brancaccio et al., 2020, 2025, 2026, Ganapati et al., 2024). Related research shows that endogenous transportation costs and market power shape trade patterns and the cost of international trade (Hummels et al., 2009, Asturias, 2020, Wong, 2022), and that transportation networks and infrastructure determine market access, the spatial allocation of activity, and welfare (Limão and Venables, 2001, Allen and Arkolakis, 2022, Fajgelbaum and Schaal, 2020, Coşar et al., 2024). Bai et al. (2026a), Fernández-Villaverde et al. (2025), and Heiland et al. (2025) also employ AIS data to trace how shocks transmit through global shipping. We differ in how we map vessel movements to route-specific measures of potential and effective capacity, and in incorporating those measures into a general equilibrium model with endogenous capacity allocation. Our model, by linking operational vessel data to equilibrium outcomes, can quantify disruptions and trace how they spill over to the rest of the network, complementing work on shipbuilding and other slower-moving changes in aggregate fleet capacity.

Second, the paper adds to research on the propagation and aggregate consequences of supply chain disruptions. A large literature studies how shocks transmit through input–output networks, firm-to-firm relationships, and international sourcing (Acemoglu et al., 2012, Boehm et al., 2019, Carvalho et al., 2021, Blaum et al., 2026). Work motivated by the 2020–2022 supply-chain crisis documents its effects on trade, production, and inflation (Alessandria et al., 2023, Bai et al., 2026a,b, Carrière-Swallow et al., 2023, Comin et al., 2023, di Giovanni et al., 2023, Isaacson and Rubinton, 2023, LaBelle and Santacreu, 2022), while related studies examine the disruption of shipping through the Red Sea (Notteboom et al., 2024, Yap and Yang, 2024) and sailing-speed adjustments (Lugovskyy et al., 2025). As we explained above, these frameworks typically represent transportation disruptions as exogenous changes in iceberg costs. Instead, in our framework, a disruption is shaped by shipping firms' capacity reallocation and importers' sourcing responses, which affect congestion and transmit the shock to prices and trade flows on every other route.

Methodologically, our analysis extends Hulten (1978) and Baqaee and Farhi (2019, 2020) to a congestible shipping network. In frictionless competitive environments, the effect of a productivity shock is summarized by the affected producer's Domar weight, and in canonical quantitative trade models, welfare changes can be recovered from domestic expenditure shares and a constant trade elasticity (Arkolakis et al., 2012). In our model, route-level trade costs depend on the endogenous allocation of a common fleet. Thus, we can show how the welfare effect of a disruption diverges both from the route's Domar weight and from the conventional domestic-expenditure sufficient statistic.

Third, our policy applications are related to research on maritime security and trade policy. The Navy application builds on work on the economic consequences of conflict and insecure trade routes (Anderson and Marcouiller, 2002, Glick and Taylor, 2010, Besley et al., 2015). This literature typically estimates the realized costs of insecurity. In our case, we use the model to value counterfactual reductions in route disruptions, estimate the benefits of naval protection in the Red Sea, and compare them with its fiscal cost.

The empirical literature on the 2018–2019 U.S. tariffs documents significant pass-through to import prices and welfare losses (Amiti et al., 2019, Cavallo et al., 2021, Fajgelbaum et al., 2020), and quantitative trade models evaluate the welfare effects of tariffs (Caliendo and Parro, 2015, Caliendo et al., 2019). Our novel contribution is to show that, when routes are congested, tariffs can act as a second-best substitute for Pigouvian congestion pricing (Bhagwati and Ramaswami, 1963, Brueckner, 2002, De Borger et al., 2011, Strandenes and Marlow, 2000).

2. U.S.-BOUND POTENTIAL SHIPPING CAPACITY

Our first task is to quantify the evolution of U.S.-bound potential shipping capacity at the route level using high-frequency satellite data. Potential capacity is the upper bound on the shipping services a route could deliver absent disruptions, measured in metric tons or ton-miles.[1]

Data. Our analysis relies on two primary data sources. The first is the Automatic Identification System (AIS), which the International Maritime Organization (IMO) requires on all ships of 300 gross tonnage and above engaged in international voyages. AIS reports the IMO ship identification number, position, and speed at high frequency. Vessels transmit as often as every two seconds. We focus on container ships, as they are the main mode of transport for man-

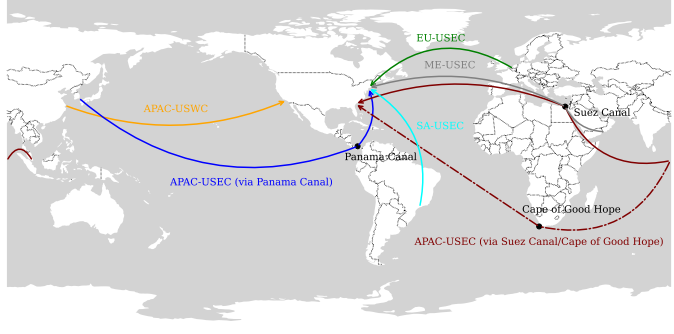
[1] A ton-mile is one metric ton of goods carried over one nautical mile (Stopford, 2008).

ufactured goods.[2] The second data source is the Lloyd's List vessel registry, which provides ship-level characteristics, including deadweight tonnage (DWT) and ownership.[3]

Combining these sources lets us reconstruct the full population of container-ship voyages arriving at U.S. ports, linked to ship-level characteristics, from January 2016, when AIS data became publicly available with comprehensive global coverage, through March 2025. In our benchmark analysis, we use weekly data unless otherwise stated.[4]

Routes and trips. We classify vessel trips into six major U.S.-bound trade routes (see Figure 2.1): (1) Asia-Pacific–West Coast (APAC–USWC); (2) Asia-Pacific–East Coast via the Panama Canal (APAC–USEC via Panama); (3) Asia-Pacific–East Coast via the Suez Canal or the Cape of Good Hope (APAC–USEC via Suez/Cape); (4) Europe–East Coast (EU–USEC); (5) Middle East–East Coast (ME–USEC); and (6) South America–East Coast (SA–USEC).[5] Together, they account for about 80–90% of U.S. containerized maritime imports.

FIGURE 2.1.—Major U.S.-bound trade routes



Unlike dry bulk vessels or tankers, container ships operate on fixed service loops connecting predetermined sets of origin and destination ports. Figure 2.2 shows a representative round trip between Asia-Pacific and the U.S. East Coast via the Suez Canal. To distinguish U.S. inbound from outbound trips, we define a trip as the journey that begins when a vessel departs its last port in the foreign region of origin (Asia-Pacific, Europe, the Middle East, or South America)

[2]Dry bulk shipping carries bulk commodities such as iron ore, coal, and grain, and oil tankers carry oil. In 2024 these segments held 43% and 28% of world fleet capacity, respectively; container ships and other general cargo ships made up the rest (UNCTAD, 2024). In the U.S., container shipping carries more tonnage (nearly one billion short tons) and value (more than $0.7 trillion) than any other transport mode, accounting for over 50% of U.S. trade by weight and roughly 30% by value (Bureau of Transportation Statistics, 2021). Some high-value items, such as computer chips, go by air, but these products depend on other components, such as motherboards, that travel by container ship.

[3]Dunn and Leibovici (2023) measure capacity in TEU capacity, a volume rather than weight measure.

[4]A key limitation of AIS-based shipping data is the inability to track cargo origins and destinations, which prevents precise measurement of cargo flows on individual routes. Integrating AIS records with U.S. bills of lading offers a promising way to overcome this limitation, and we leave it for future research.

[5]For route (3), we combine the Suez Canal and Cape of Good Hope paths because the latter emerged in response to the Houthi attacks in the Red Sea in late 2023 and is not a regular route throughout the period.

and ends on its arrival at the first port on either the U.S. East Coast or West Coast, possibly after some intermediate port calls (a port call is the arrival of a ship at a port where it docks to load and unload cargo). When a vessel's itinerary includes stops in multiple foreign regions, we record a separate trip for each region. This granularity lets us focus on shipping dynamics along the intercontinental legs of vessels' service loops, while limiting the effect of domestic and within-region movements on the measurement of shipping services provided to the U.S.[6]

FIGURE 2.2.—Representative service loop: Asia-Pacific to U.S. East Coast via Suez



Note: A representative port rotation (U.S.-inbound leg in blue, U.S.-outbound leg in gray) for a container service connecting the Asia-Pacific and Middle East regions to the U.S. East Coast via the Suez Canal. The route follows a fixed sequence of port calls, from East Asia (Qingdao, Shanghai, Ningbo, and Nansha) to U.S. East Coast ports (Newark, Norfolk, Charleston, and Jacksonville). The vessel then returns to Asia to complete the round trip. The schedule on the left reports the pro forma departure day at each port along the service loop. *Source:* SeaLead.

Table 2.1 summarizes the main destination and origin ports in our sample, along with the routes they follow and the share of trips they represent over 2016–2025. The sample records 58,582 trips and includes 218 unique ports appearing as either destinations or foreign-origin ports, comprising 58 destination ports and 160 origin ports. The principal U.S. destination ports are Newark on the East Coast, Los Angeles and Long Beach on the West Coast, and Savannah in the Southeast. The leading Asia-Pacific origin ports are Busan in South Korea and Shanghai, Yantian, and Ningbo-Zhoushan in China, primarily associated with the APAC–USWC and APAC–USEC routes. Cartagena, Colombia, is the dominant South American origin port on the SA–USEC route. The main European origin ports are Valencia in Spain, Bremerhaven in Germany, and Le Havre in France, which are associated with the EU–USEC route.

Potential shipping capacity. The potential shipping capacity in *metric tons* along route r in week t is the total DWT of the container ships operating on it, where a vessel's DWT is the maximum weight of cargo it can carry:

$$\texttt{Potential Capacity}^{\text{tons}}_{r,t} = \sum_{i \in \boldsymbol{\mathcal{V}}_t} \mathbf{1}_{i,r,t} \times \texttt{DWT}_i, \tag{2.1}$$

where $\boldsymbol{\mathcal{V}}_t$ denotes the set of active container ships in week t (vessels recording at least one AIS signal), and $\mathbf{1}_{i,r,t}$ is equal to one if ship i operates on route r in week t and zero otherwise.

[6]The analysis can be extended to within-region legs or additional U.S.-bound routes, but doing so requires additional choices on route definitions and the treatment of intermediate movements, which we leave for future work.

TABLE 2.1

KEY PORTS AND U.S.-BOUND TRADE ROUTES

Rank	Destination Port	Trips (%)	Rank	Origin Port	Trips (%)	Major Routes
1	Newark, NJ	19.6	1	Busan, South Korea	19.3	APAC-USWC, APAC-USEC (Panama)
2	Los Angeles, CA	11.4	2	Cartagena, Colombia	9.8	SA-USEC
3	Long Beach, CA	9.0	3	Shanghai, China	6.8	APAC-USWC, APAC-USEC (Panama)
4	Savannah, GA	5.3	4	Yantian, China	5.5	APAC-USWC, APAC-USEC (Panama)
5	Fort Lauderdale, FL	4.3	5	Ningbo-Zhoushan, China	4.8	APAC-USWC, APAC-USEC (Panama)
6	Seattle, WA	4.0	6	Valencia, Spain	4.6	EU-USEC
7	Tacoma, WA	3.9	7	Bremerhaven, Germany	3.5	EU-USEC
8	Houston, TX	3.8	8	Le Havre, France	3.0	EU-USEC
9	Miami, FL	3.6	9	Colombo, Sri Lanka	2.6	APAC-USEC (Suez/Cape)
10	Philadelphia, PA	2.7	10	Tokyo, Japan	2.4	APAC-USWC, APAC-USEC (Panama)
11	Charleston, SC	2.7	11	Yokohama, Japan	2.2	APAC-USWC
12	New York, NY	1.6	12	Kaohsiung, Taiwan	1.9	APAC-USWC, APAC-USEC (Panama)
13	Norfolk, VA	1.4	13	Santa Marta, Colombia	1.8	SA-USEC
14	Oakland, CA	1.3	14	Singapore, Singapore	1.7	APAC-USEC (Suez/Cape)
15	Portsmouth, VA	1.2	15	Point Lisas, Trinidad & Tobago	1.5	SA-USEC

Note: The third and sixth columns show the percentage of total container-ship trips accounted for by each destination and origin port over the sample period 2016–2025. The last column lists the major routes served from each origin port.

The potential shipping capacity in *ton-miles* is the product of total DWT and the potential distance vessels could travel in one week, assuming they operate at the maximum speed observed over the sample period $\mathcal{T}$. We use the maximum observed speed because it serves as a technological upper bound independent of transient operational slowdowns. Let $\overline{\texttt{Speed}}_{r,t}$ denote the DWT-weighted average vessel speed along route r in week t:

$$\overline{\texttt{Speed}}_{r,t} = \sum_{i \in \mathcal{V}_t} \left(\frac{\mathbf{1}_{i,r,t} \times \texttt{DWT}_i}{\sum_{i' \in \mathcal{V}_t} \mathbf{1}_{i',r,t} \times \texttt{DWT}_{i'}} \right) \times \texttt{Speed}_{i,t}. \tag{2.2}$$

Since the distance traveled is the product of speed and time, the weekly potential shipping capacity in *ton-miles* is:

$$\texttt{Potential Capacity}^{\text{ton-miles}}_{r,t} = \texttt{Potential Capacity}^{\text{tons}}_{r,t} \times \underbrace{\left(\max_{s \in \mathcal{T}} \left\{ \overline{\texttt{Speed}}_{r,s} \right\} \times 24\,\text{hours} \times 7\,\text{days} \right)}_{\text{potential distance}}. \tag{2.3}$$

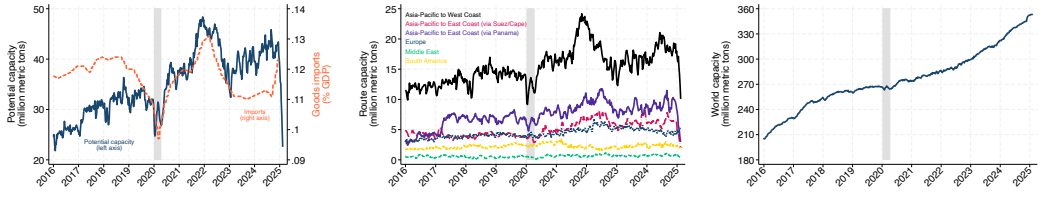
This metric captures not only the weight of the goods transported, but also the distance they are carried. Moving 10,000 tons of cargo 5,000 nautical miles demands far more shipping capacity than moving the same cargo 500 nautical miles.

The supply of U.S.-bound capacity. Panel 2.3a plots the aggregate potential shipping capacity in tons, summed across all routes.[7] From early 2016 onward, the total DWT of U.S.-bound ships rose steadily, averaging about 30 million metric tons just before the COVID-19 pandemic. After the onset of the pandemic, potential capacity surged by more than 60%, reaching 48 million metric tons in January 2022. This expansion reflected both the entry of new firms

[7] Aggregating across routes may overstate potential capacity because a single vessel can serve multiple routes. Figure A1 of Appendix A addresses this with the set of unique container ships on U.S.-bound routes; the magnitudes differ slightly but the dynamics are essentially the same as in Figure 2.3.

and the redeployment of vessels from other regions to capture historically high freight rates on American routes. Total DWT then contracted from its peak as import demand normalized, before rebounding toward the end of 2024. In early 2025, potential shipping capacity fell sharply, ahead of the escalation of the U.S.–China trade war.[8]

FIGURE 2.3.—U.S.-bound potential shipping capacity (tons), January 2016–March 2025



(a) U.S.-bound potential capacity (b) Route-level potential capacity (c) World potential capacity

Note: Panel 2.3a shows aggregate U.S.-bound potential capacity in million metric tons (left axis) and U.S. goods imports as a share of GDP (right axis). The shaded area indicates recessions (this note also applies to the figures in Section 3).

Panel 2.3b shows that the expansion of shipping capacity toward the U.S. was uneven across routes. The largest increase occurred on the APAC–USWC route.By the end of 2021, this route accounted for roughly half of total U.S.-bound potential capacity in tons, up about 6 percentage points from early 2020.

Finally, Panel 2.3c shows that the world's total capacity evolved much more steadily, suggesting that reallocation occurred mostly across regions. Africa (12.4%) and Latin America and the Caribbean (13.5%) lost the largest shares of their direct liner shipping connections to this redeployment between late 2020 and mid-2022 (UNCTAD, 2022). Carriers preferred the APAC-USWC route because before the pandemic shipments took only 2.2 weeks on average, against about 4 weeks from Asia-Pacific to the East Coast (shipments from Europe, the Middle East, and South America averaged 1.7, 2.9, and 1.3 weeks). Shorter trips raised fleet turnover, pulling capacity to the West Coast when import demand peaked.

Although less visible in Panel 2.3b, potential shipping capacity from the APAC–USEC routes (via either Panama or Suez/Cape) also expanded gradually, peaking toward the end of 2022. The delayed adjustment reflects that rerouting to the East Coast required longer sailing distances, higher fuel costs, and more time to reposition fleets. Hence, carriers redeployed capacity first on the more profitable APAC–USWC route (Notteboom et al., 2021). Similar patterns emerged on the EU–USEC and ME–USEC routes. By contrast, the capacity of the SA–USEC route

[8]The series in ton-miles (Panel 3.5a) traces the same path, from a 2016–2019 average of slightly more than 80 billion to a peak of 135 billion in early 2022, about 110 billion in 2023–2024, and below 70 billion by the end of the sample. The number of container ships serving U.S. ports moved through the same cycle (Figure A2 of Appendix A), rising from about 360 a week before the pandemic to 550 in early 2022 before correcting in 2023.

declined substantially from early 2021 to mid-2022, as shipping lines reassigned vessels to the Asia–U.S. routes and away from Latin America and the Caribbean.

3. U.S.-BOUND EFFECTIVE SHIPPING CAPACITY AND UTILIZATION

Not all the surge in potential shipping capacity after 2020 translated into effective shipping services. Port congestion, bottlenecks at the Suez and Panama Canals, rerouting around conflict zones, and slow steaming constrained the shipping supply to U.S. importers, both by stalling vessels at ports or canals and by forcing ships to sail farther or slower.

To measure the actual provision of shipping services, we compute effective capacity in ton-miles, the capacity that remains after accounting for operational frictions, or

$$\begin{aligned}\texttt{Effective Capacity}_{r,t}^{\text{ton-miles}} = \texttt{Potential Capacity}_{r,t}^{\text{ton-miles}} - \texttt{Port Congestion}_{r,t}^{\text{ton-miles}} \\ - \texttt{Canal Congestion}_{r,t}^{\text{ton-miles}} - \texttt{Rerouting}_{r,t}^{\text{ton-miles}} \\ - \texttt{Vessel Slowdown}_{r,t}^{\text{ton-miles}}.\end{aligned} \tag{3.1}$$

Our measure of shipping utilization is the ratio of effective to potential shipping capacity:

$$\texttt{Shipping Utilization}_{r,t} = \texttt{Effective Capacity}_{r,t}^{\text{ton-miles}} / \texttt{Potential Capacity}_{r,t}^{\text{ton-miles}}. \tag{3.2}$$

Shipping utilization summarizes the impact of disruptions in a single unit-free statistic, comparable across routes and over time (unlike days of delay or queue lengths).

Port congestion. We measure port congestion using AIS records of vessel speed and anchorage duration. Congestion occurs when ships must wait in a port's anchorage area until a berth becomes available (Bai et al., 2026a). Ideally, a ship proceeds straight to a berth on arrival. In practice, disruptions ranging from labor shortages and inland transport delays to surges in import demand force vessels to queue at anchor, reducing the actual supply of shipping services.

We quantify this loss in ton-miles per route at a weekly frequency. First, we map the geographic boundaries of berth and anchorage areas for all container ports worldwide.[9] Next, for each port call we compute the time the vessel spends in that anchorage area, attributing the full wait to the week in which the call began. Let $\texttt{Port Waiting Time}_{i,p,r,t}$ denote the hours vessel i spends in anchorage at port p during a port call initiated in week t on a trip along route r.

To convert this waiting time into foregone nautical miles, we assume that absent congestion the vessel would have sailed at the route's average speed. We therefore multiply the anchorage duration by the DWT-weighted average speed of all vessels on route r in week t ($\overline{\texttt{Speed}}_{r,t}$ in

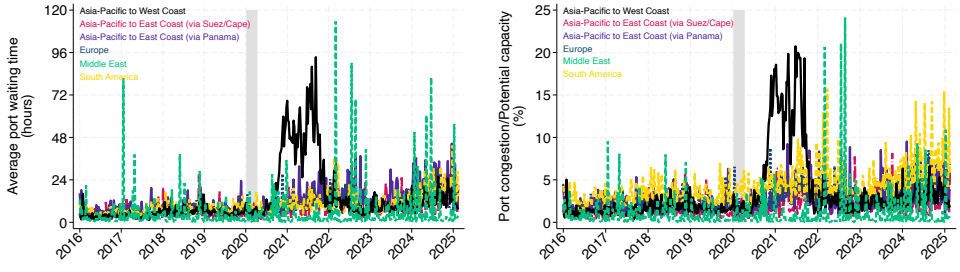
[9] We use AIS data with the IMA-DBSCAN (Iterative Multi-Attribute Density-Based Spatial Clustering of Applications with Noise) algorithm for this task; see Bai et al. (2023) for details.

Equation (2.2)) and then by the vessel's DWT, yielding:

$$\texttt{Port Congestion}_{r,t}^{\text{ton-miles}} = \sum_{i \in \boldsymbol{\mathcal{V}}_t} \mathbf{1}_{i,r,t} \times \underbrace{\texttt{DWT}_i}_{\text{tonnage}} \times \underbrace{\left(\overline{\texttt{Speed}}_{r,t} \times \sum_{p \in \boldsymbol{\mathcal{P}}_{i,r,t}} \texttt{Port Waiting Time}_{i,p,r,t} \right)}_{\text{foregone distance}}, \tag{3.3}$$

where $\boldsymbol{\mathcal{P}}_{i,r,t}$ is the set of port calls initiated by vessel i along route r in week t, and $\boldsymbol{\mathcal{V}}_t$ is the set of active container ships in week t.

FIGURE 3.1.—Port congestion across routes, January 2016–March 2025



(a) Average port waiting time

(b) Port congestion/potential capacity

Panel 3.1a shows the average waiting time at port anchorages. Across all routes, ships spent roughly 6 hours at anchor before the pandemic. Driven by a surge in consumer demand and pandemic-related labor shortages at the Ports of Los Angeles and Long Beach, average waiting times on the APAC–USWC route repeatedly exceeded 60 hours in the summer and fall of 2021. As West Coast bottlenecks worsened while national demand stayed high, shipping companies moved capacity to East Coast ports. That shift eased congestion on the APAC–USWC route but worsened it on the East Coast. Delays on the ME–USEC route were severe, with average anchorage times peaking above 80 hours in some weeks.[10]

Panel 3.1b plots port congestion as a fraction of potential capacity in ton-miles. Port congestion cut potential shipping capacity on the APAC–USWC route by as much as one-fifth in 2021, a substantial utilization loss. The effect was similarly severe on the ME–USEC route, where port congestion absorbed nearly a quarter of potential capacity in late 2022.

Canal congestion. Much of global shipping crosses the checkpoints of the Suez and Panama Canals. Thus, we track vessel movements within the canal zones, identify periods when ships are stalled or delayed, and compute the ton-miles lost to queuing at them.

[10]The ME–USEC series is volatile because a very small fleet serves the lane, fewer than 10 vessels for most of the sample period. With so few ships, the decisions of a single vessel or a delay at a regional chokepoint can move the route average a lot.

First, we define the geographic boundaries of the Suez and Panama Canals with polygons. Within these boundaries, we identify all vessel transits by extracting the sequence of AIS signals each vessel broadcasts. For every pair of consecutive signals, we compute the time elapsed, the distance traveled, and the implied speed over that interval.

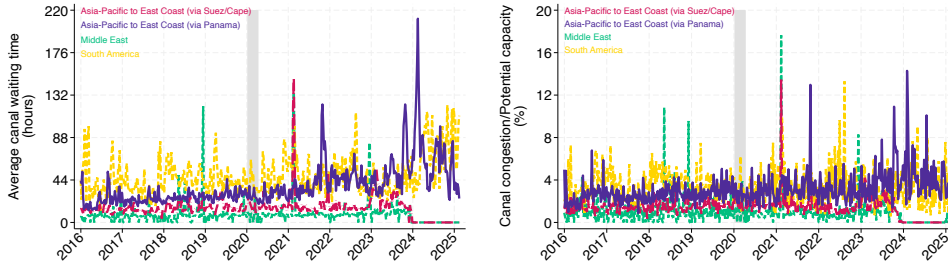
Next, we isolate intervals where a vessel is stationary or moving well below normal transit speed because of canal bottlenecks, filtering for intervals where speed is one knot or less. For these delayed intervals, we compute the foregone distance, the distance the vessel would have covered at the route's benchmark speed minus the distance it actually covered. The benchmark speed is $\overline{\texttt{Speed}}_{r,t}$ from Equation (2.2).[11]

The measure of canal congestion multiplies the foregone distance by the vessel's DWT and sums across all valid intervals for all ships on the route:

$$\texttt{Canal Congestion}_{r,t}^{\text{ton-miles}} = \sum_{i \in \mathcal{V}_t} \mathbf{1}_{i,r,t} \times \underbrace{\texttt{DWT}_i}_{\text{tonnage}} \times \underbrace{\left[\sum_{k \in \mathcal{K}_{i,r,t}} \mathbf{1}_{\texttt{Speed}_{i,k} \leq 1} \times \left(\overline{\texttt{Speed}}_{r,t} - \texttt{Speed}_{i,k} \right) \times \tau_{i,k} \right]}_{\text{foregone distance}}, \tag{3.4}$$

where $\tau_{i,k}$ is the duration in hours of each interval k between consecutive AIS signals within the boundaries of the Suez or Panama Canal for vessel i, and $\mathcal{K}_{i,r,t}$ is the set of such intervals for vessel i along route r in week t. $\texttt{Speed}_{i,k}$ is the vessel's speed in interval k, and $\mathbf{1}_{\texttt{Speed}_{i,k} \leq 1}$ ensures we aggregate foregone miles only during periods when the vessel was effectively stalled.

FIGURE 3.2.—Canal congestion across routes, January 2016–March 2025



(a) Average canal waiting time

(b) Canal congestion/potential capacity

Panel 3.2a shows the average waiting time at the Suez and Panama Canals for routes from Asia-Pacific, the Middle East, and South America to the U.S. East Coast.[12] Before the pan-

[11] In Equation (3.3), we do not subtract the distance traveled, since vessels in anchorage are mostly stationary at zero speed. In the canal congestion calculation we account for the fact that waiting vessels may still move slowly through the canal in queues.

[12] These chokepoints do not matter for the APAC–USWC and EU–USEC routes, which rely on direct trans-Pacific or trans-Atlantic connections.

demic, waiting times were modest: 11 hours at Suez and 34 hours at Panama. Since then, the two canals have diverged. Suez had a short-lived spike in March 2021, when the *Ever Given* blocked the canal, and another in January 2023, when the bulk carrier *M/V Glory* ran aground but was quickly refloated. Since early 2024, unrest in the Red Sea has pushed many carriers onto the longer route around the Cape of Good Hope (see Figure 3.3), cutting Suez traffic and waiting times to near zero. In comparison, waiting times at the Panama Canal have risen since 2023 due to severe drought and the onset of El Niño cut transit capacity.

Panel 3.2b plots canal congestion as a fraction of potential capacity across routes. Suez congestion accounted for less than 2% of potential capacity on the APAC–USEC (via Suez/Cape) and ME–USEC routes for most of the sample, then fell to zero as the Red Sea crisis escalated. The 2023–2024 drought and El Niño affected vessels crossing the Panama Canal, absorbing more than 4% of potential capacity on the APAC–USEC (via Panama) and SA–USEC routes.

Rerouting. Next, we examine rerouting episodes such as the 2023–2024 Red Sea crisis. In this episodes, vessels are forced off their planned course onto a longer or costlier alternatives. We start by grouping container-ship trips along each U.S.-bound route by their origin–destination (o–d) port pairs. We focus on routes that offer multiple options (e.g., Suez and the Cape of Good Hope). This ensures the measured deviations reflect material changes in routing rather than minor navigational adjustments on a single fixed route.

Since, under normal conditions, vessels traverse the shortest available path to minimize fuel consumption and transit time, we identify the trip with the shortest travel distance and use the corresponding vessel's trajectory for each o–d pair and label it our benchmark. Let $\texttt{Min Distance}_{o,d,r}$ denote the shortest trip distance between origin o and destination d along route r, and let $\texttt{Trip Distance}_{i,j,o,d,r}$ denote the distance for vessel i's trip j from o to d along route r, in nautical miles.

We classify a trip j as rerouted if it deviates from the benchmark trajectory. Since trips span varying durations, we measure the intensity of the deviation as a rate. The additional nautical miles of travel in a week are

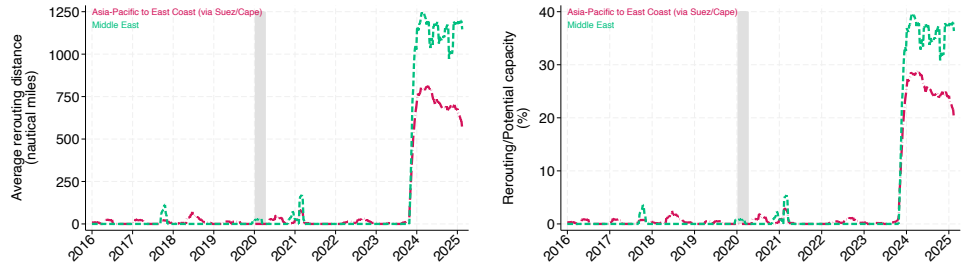
$$\texttt{Weekly Extra NM}_{i,j,o,d,r} \equiv \underbrace{\frac{\max\left(0, \texttt{Trip Distance}_{i,j,o,d,r} - \texttt{Min Distance}_{o,d,r}\right)}{\texttt{Trip Duration}_{i,j,o,d,r}}}_{\text{extra nautical miles per hour}} \times 24\,\text{hours} \times 7\,\text{days},$$

where $\texttt{Trip Duration}_{i,j,o,d,r}$ is trip duration in hours. The rerouting measure in ton-miles is

$$\texttt{Rerouting}^{\text{ton-miles}}_{r,t} = \sum_{o \in \boldsymbol{\mathcal{O}}_r} \sum_{d \in \boldsymbol{\mathcal{D}}_r} \sum_{j \in \boldsymbol{\mathcal{J}}_{o,d,r,t}} \sum_{i \in \boldsymbol{\mathcal{V}}_t} \mathbf{1}_{i,j} \times \underbrace{\texttt{DWT}_i}_{\text{tonnage}} \times \underbrace{\texttt{Weekly Extra NM}_{i,j,o,d,r}}_{\text{foregone distance}}, \tag{3.5}$$

where $\boldsymbol{\mathcal{O}}_r$ and $\boldsymbol{\mathcal{D}}_r$ are the sets of origin and destination ports along route r, $\boldsymbol{\mathcal{J}}_{o,d,r,t}$ is the set of rerouted trips between o and d along route r in week t, and $\boldsymbol{\mathcal{V}}_t$ is the set of active vessels in week t. The indicator $\mathbf{1}_{i,j}$ equals one if vessel i performed trip j and zero otherwise.

FIGURE 3.3.—Vessel rerouting, January 2016–March 2025



(a) Average rerouting distance

(b) Rerouting/potential capacity

Panel 3.3a shows the extra distance, in nautical miles, sailed by vessels rerouting from lanes that normally transit the Suez Canal. The March 2021 Suez blockage caused negligible deviations, but the Red Sea crisis that began in late 2023 brought a lasting changes. Going around the Cape of Good Hope added roughly 600–800 nautical miles to voyages from the Middle East and 1,000–1,250 to voyages from Asia-Pacific to the U.S. East Coast. Panel 3.3b shows that in 2024 rerouting absorbed about 25% of the potential capacity deployed on the APAC–USEC (via Suez/Cape) route and 36% of that on the ME–USEC route.

Vessel slowdown. We account for reductions in shipping capacity from speed adjustments. Even when diverting to longer routes, vessels may operate below their usual speed to save fuel or manage arrival times at congested ports, a practice known as *slow steaming*.

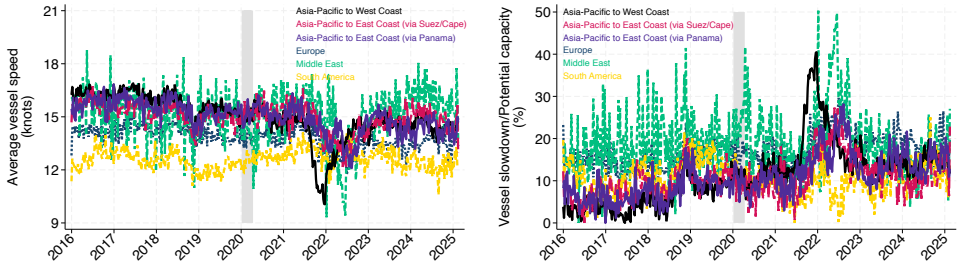
Vessel slowdown is the gap between the average fleet speed ($\overline{\texttt{Speed}}_{r,t}$ from Equation (2.2)) and the route's technological potential, the maximum average speed observed on route r over the sample period $\mathcal{T}$. The capacity lost to vessel slowdown is the product of aggregate deployed tonnage and the distance foregone per week by sailing below this maximum:

$$\texttt{Vessel Slowdown}^{\text{ton-miles}}_{r,t} = \left[\sum_{i \in \mathcal{V}_t} \mathbf{1}_{i,r,t} \times \underbrace{\texttt{DWT}_i}_{\text{tonnage}}\right] \times \underbrace{\left[\left(\max_{s \in \mathcal{T}} \left\{\overline{\texttt{Speed}}_{r,s}\right\} - \overline{\texttt{Speed}}_{r,t}\right) \times 24\,\text{hours} \times 7\,\text{days}\right]}_{\text{foregone distance}}. \tag{3.6}$$

Panel 3.4a plots the evolution of average vessel speed across the six U.S.-bound routes. Before the pandemic (2016–2019), vessels on the APAC–USWC route averaged about 16 knots but speeds dropped sharply during the pandemic. Around the turn of 2022, speeds were just near 10 knots. The deceleration was due largely to a new virtual queuing system introduced at the Ports of Los Angeles and Long Beach in late 2021.[13] The system, designed to reduce safety risks and pollution from ships waiting near port, assigned queue positions by departure time,

[13] See MXSOCAL (2021) (accessed Dec 17, 2025) for details.

FIGURE 3.4.—Vessel slowdown across routes, January 2016–March 2025



(a) Average vessel speed

(b) Vessel slowdown/potential capacity

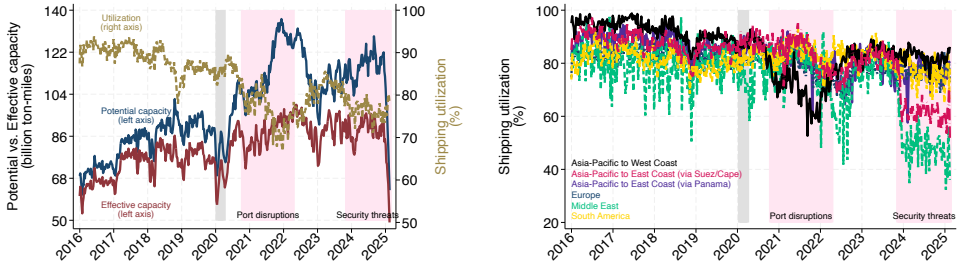
removing the "hurry up and wait" incentive and encouraging vessels to slow steam across the Pacific.

Average speeds on the EU–USEC and APAC–USEC (via Panama) routes fell from 14.2 and 15.5 knots before the pandemic to about 13.5 and 14.5 knots. High fuel prices in 2022 and environmental rules, such as the IMO's Carbon Intensity Indicator (which forced many older vessels to sail more slowly to meet carbon intensity targets), kept speeds down. Speeds have picked up a little since mid-2023 but remain below pre-pandemic levels.

Panel 3.4b reports the share of potential capacity absorbed by vessel slowdown. On the APAC–USWC route, slow steaming accounted for about 6% of potential capacity before the pandemic, but nearly 40% at the height of the supply chain crisis. The loss is not unique to the Pacific: on the EU–USEC route the share rose from 15% pre-pandemic to 20% more recently.

The shipping utilization collapse. Panel 3.5a presents potential versus effective shipping capacity (in ton-miles) aggregated across all U.S.-bound routes. The gap between the two is the total capacity lost to shipping disruptions. The panel also reports aggregate shipping utilization.

FIGURE 3.5.—U.S.-bound effective capacity and shipping utilization, January 2016–March 2025



(a) Aggregate trends

(b) Shipping utilization

Figure 3.5 shows that, in 2016–2019, the U.S.-bound shipping network turned about 89% of the fleet's potential capacity into effective service; normal operational frictions absorbed the other 11%. During 2021–2022, effective capacity declined despite potential capacity reaching record highs. At the worst of the disruption, aggregate utilization of the U.S.-bound fleet fell below 70% due to ships waiting at anchor, steaming slowly, or sailing around maritime bottlenecks. The network has not regained its pre-pandemic utilization. More recently, rerouting after the 2023–2024 Houthi attacks in the Red Sea held aggregate utilization near 79%.

The aggregate decline hides wide heterogeneity across routes. Panel 3.5b plots shipping utilization for the six U.S.-bound routes. The APAC–USWC route sat at the center of the pandemic-era collapse. While before 2020, this route had been the most efficient corridor, turning 92% of potential capacity into effective supply, by late 2021, port congestion and slower steaming pushed utilization down to 53%. The route has since recovered but remains well below its old peak.

The recent deterioration, by contrast, owes largely to geopolitical instability following the Houthi attacks in the Red Sea. Utilization on the Middle East and Asia-Pacific routes to the U.S. East Coast via the Suez Canal has fallen sharply since late 2023, to 48% and 61%, respectively. Diverting vessels around the Cape of Good Hope drives much of this loss, since ships travel longer distances without delivering additional cargo.

4. A MODEL OF THE SHIPPING MARKET WITH MULTIPLE ROUTES

The evidence in Sections 2 and 3 points to two complementary forces in U.S.-bound shipping. Shipping firms allocate and reallocate capacity across routes, while disruptions from congestion and security threats absorb large shares of that capacity, cutting the effective supply of shipping services reaching U.S. ports.

To quantify these patterns, we develop a multi-route general equilibrium model with both forces: (i) the endogenous allocation of capacity across routes by profit-maximizing firms, and (ii) disruptions that govern how much potential capacity becomes effective shipping services. The model lets us trace how shocks propagate from the shipping sector to trade flows and prices, and evaluate counterfactuals central to trade policy.

4.1. *Environment*

Consider an economy with consumers, producers, retailers, and shipping firms. Besides domestically produced goods, consumers enjoy foreign goods that must be shipped from one of the four origin regions $o \in \mathcal{O} = \{A, E, M, S\}$ (Asia, Europe, the Middle East, and South America) to one of the two U.S. arrival ports $a \in \mathcal{A} = \{E, W\}$ (the East and West Coast). Imported goods from different regions are imperfect substitutes.

Shipping firms operate a fleet of container vessels, competing across multiple U.S.-bound routes. They carry the manufactured goods produced in one of the four regions to a destination port along a route $r \in \mathcal{R}$. The six routes in $\mathcal{R}$ are *Asia to West Coast*, *Asia to East Coast via Panama Canal*, *Asia to East Coast via Suez Canal*, *Europe to East Coast*, *Middle East to East Coast*, and *South America to East Coast*.[14] A shipping firm first decides how to allocate vessel capacity across routes, subject to a capacity constraint and costs of reshuffling fleets, repurposing vessels, and repositioning crews. Given its optimal capacity allocation, the firm then makes the operational decisions of which sailing speeds to adopt on each route and how intensively to load ships. Once ships arrive, importers resell the goods to consumers.

Domestic goods producers. A representative firm produces domestic goods d with a linear technology in labor with productivity z_d. The firm hires domestic labor, l_d, at the wage rate w and sells the goods directly to consumers at price p_d. Profit maximization requires:

$$p_d z_d = w. \tag{4.1}$$

Foreign goods producers. Foreign firms produce goods for export in each origin region. They supply any quantity demanded at the price $\overline{p}_o$ (in dollars) for goods originating in region $o \in \mathcal{O}$. Foreign producers sell to importers, but for simplicity, we have the shipping firms act as intermediaries. The producers are indifferent to the routes shipping firms take and so do not price discriminate across them (relaxing this assumption is straightforward).

Shipping market. There is a continuum of monopolistically competitive shipping firms of measure N. Each shipping firm is endowed with total capacity κ that can be optimally allocated among different routes or remain idle. The capacity allocated to each route r is denoted by κ_r and idle capacity (or capacity allocated to unmodeled routes) is $\overline{\kappa}$.[15] The capacity constraint of a shipping firm can be expressed as

$$\kappa = \sum_{r \in \mathcal{R}} \kappa_r + \overline{\kappa}. \tag{4.2}$$

Shipping technology. Along each route r, shipping firms choose the average speed of their ships, s, and their loading intensity, ℓ, to produce the shipping service q_r. The technology to produce the shipping service is a function of the number of trips made along the route, $n_r(s)$,

[14] The six routes correspond to the empirical routes APAC–USWC, APAC–USEC (via Panama), APAC–USEC (via Suez/Cape), EU–USEC, ME–USEC, and SA–USEC of Sections 2 and 3.

[15] As new ships take on average three years to build (Kalouptsidi, 2014), idle capacity plays a role in counterfactual experiments as it puts an upper bound on how much capacity shipping firms can reallocate to U.S.-bound routes.

and the quantity of goods transported per trip, $\ell\kappa_r$:

$$\underbrace{q_r}_{\text{effective supply}} = z_r \underbrace{n_r(s)\,\ell\kappa_r}_{\text{potential supply}}. \tag{4.3}$$

Here z_r is a firm- and route-specific capacity utilization term that maps to the shipping capacity utilization measured in Section 3. It is in turn given by

$$z_r = \frac{x_r}{g(K_r)}, \tag{4.4}$$

where x_r is an exogenous route-specific disruption component capturing frictions outside the control of shipping firms, such as labor shortages, port and canal blockages, security threats, and extreme weather, and $g(K_r)$ is a congestion externality that rises with the market capacity allocated to the route, $K_r \equiv \int_0^N \kappa_r(j)\,dj$. A fall in z_r denotes an increase in disruptions.

The number of trips made within a period is

$$n_r(s) = A_r s^{\alpha}, \tag{4.5}$$

where A_r is a route-specific efficiency term translating nautical miles per hour into number of trips, and $\alpha \in (0,1)$ is a parameter.[16] The average speed of ships is the ratio of the distance traveled along route r, δ_r, to the time traveled within the period, τ, that is, $s \equiv \delta_r/\tau$.[17]

Finally, the loading of ships, $\ell \in [0,1]$, is the fraction of the capacity allocated to the route, κ_r, used to transport goods. If $\ell = 0$, ships travel empty; if $\ell = 1$, they are at full capacity.

Congestion. The congestion factor disrupting route r depends on the aggregate capacity deployed along the route by all competing shipping firms:

$$g(K_r) = \varepsilon_k \left(\int_0^N \kappa_r(j) dj \right)^{\phi}, \tag{4.6}$$

where ε_k is a shifter and ϕ is the elasticity of the congestion externality with respect to the total capacity traveling along route r. The more capacity (i.e., container ships) travels along a route, the more congested and disrupted it is.

Shipping costs. The transportation cost per trip is $\overline{p}_n$, assumed common across routes. The more trips a shipping firm makes, the costlier it is to transport goods. Because shipping firms act as intermediaries that purchase the goods they carry, shipping costs also include the expenditure

[16] Since each trip includes port time that does not scale with speed, trips rise less than proportionally with speed.

[17] The distance along route r is exogenous and common to all shipping firms, so increases in speed reduce the time traveled between origin and arrival ports.

on goods at the origin price $\overline{p}_o$. The total cost of shipping along route r is

$$tc_r = \overline{p}_n (n_r(s))^{\zeta} + \overline{p}_o (\ell \kappa_r)^{\nu}. \tag{4.7}$$

The cost exponents are such that $\zeta, \nu \geq 1$ and $1/\zeta + 1/\nu = 1$. This restriction ensures the total cost function is homogeneous of degree 1 in shipping output.[18] The total cost is convex in trips and in goods transported, reflecting diminishing returns to, or rising adjustment costs from, faster sailing (more engine maintenance, more specialized sailors) or heavier loading (the cranes and workers needed to stack containers grow costlier).

Operational decisions. The shipping firm's problem has a two-stage structure. The firm first commits to a cross-route capacity allocation, then chooses speed and loading on each route to maximize per-route profits. The optimal sailing speed and loading, s_r and ℓ_r, minimize the firm's cost from operating in route r, Equation (4.7), subject to the technological constraint, Equation (4.3), and the loading constraint $\ell \in [0, 1]$, and taking the capacity allocated to the route, κ_r, as given. When the optimal load is interior, that is, $\ell_r < 1$, the solution yields the marginal cost of shipping through route r:

$$mc_r = \frac{g(K_r)}{x_r} \left[\frac{\overline{p}_n}{1/\zeta}\right]^{1/\zeta} \left[\frac{\overline{p}_o}{1/\nu}\right]^{1/\nu}. \tag{4.8}$$

The marginal cost of shipping rises with disruptions, congestion, and the prices of transportation and of goods at the origin port. When the ship is fully loaded, it also rises with capacity.[19]

Shipping profits per route. A shipping firm maximizes its profits along route r by choosing the price of its shipping services, p, which bundles the cost of the goods with that of the voyage, taking the importers' demand $\mathtt{y}_r(p)$ as given. The route-r profit maximization problem is

$$\pi_r = \max_p \left(p - mc_r\right) q_r, \tag{4.9}$$

subject to supply meeting demand, $q_r = \mathtt{y}_r(p)$.

The solution yields a route-specific price, p_r, equal to a markup over marginal cost:

$$p_r = m_r \, mc_r. \tag{4.10}$$

[18] The exponents $1/\zeta$ and $1/\nu$ in the cost function are isomorphic to having ζ and ν in the production technology.

[19] When the optimal load is at the boundary $\ell_r = 1$, the marginal cost of shipping through route r is $mc_r = \frac{g(K_r)}{x_r} \left[\frac{\overline{p}_n}{1/\zeta}\right]^{1/\zeta} \left[\frac{\overline{p}_o}{1/\nu} + \frac{\mu_r}{\kappa_r^{\nu}}\right]^{1/\nu}$, where $\mu_r > 0$ is the Lagrange multiplier on the load's upper-bound constraint. The multiplier can be written as a function of capacity, number of trips, and input prices: $\mu_r = \zeta \overline{p}_n n_r (s_r)^{\zeta} - \nu \overline{p}_o \kappa_r^{\nu}$.

The importers' price elasticity of demand, $\xi_r \equiv -(\partial \mathtt{y}_r(p_r)/\partial p_r)(p_r/\mathtt{y}_r(p_r))$, determines the markup through $m_r = \xi_r/(\xi_r - 1)$. In the decentralized equilibrium, the demand in Equation (4.17) has a constant elasticity $\xi_r = 1/(1-\sigma)$, yielding the same markup across all routes, $m_r = 1/\sigma$.

Optimal capacity per route. The capacity allocated to each route r is chosen to maximize the shipping firm's total profit, π, the sum of per-route profits net of the cost of deploying ships across routes. That cost is a quadratic capacity-adjustment term $\overline{p}_\kappa$, which captures the expense of repositioning ships, repurposing vessel types, and reshuffling crews across routes. The shipping firm's total profit is

$$\pi = \max_{\{\kappa_r\}_{r\in\mathcal{R}}} \left\{ \sum_r \pi_r - \overline{p}_\kappa \sum_r \frac{\kappa_r^2}{2} \right\}, \tag{4.11}$$

subject to the capacity constraint (4.2). The optimal capacity per route follows from substituting the Lagrange multiplier on the capacity constraint into the firm's first-order condition for problem (4.11). The firm does not internalize the effect of its capacity choice on congestion, and hence on output.

The optimal capacity per route equates the marginal benefit of additional capacity to its marginal cost. Adding capacity on one route raises that route's profit per ton traveled relative to the firm's average, but it also raises the cost of deploying capacity there relative to the average. The more congested a route, the lower its profit per ton traveled and the more likely the firm is to reallocate capacity to other routes. The first-order condition, when the optimal loading is interior ($\ell_r < 1$), is

$$\underbrace{\frac{(p_r - mc_r)\, x_r n_r(s_r^\star)\ell_r^\star}{g(K_r)}}_{\text{profit per ton traveled along route } r} - \underbrace{\frac{1}{|\mathcal{R}|}\sum_h \frac{(p_h - mc_h)\, x_h n_h(s_h^\star)\ell_h^\star}{g(K_h)}}_{\text{average profit per ton traveled}} = \overline{p}_\kappa \left[\kappa_r - \underbrace{\frac{1}{|\mathcal{R}|}(\kappa - \overline{\kappa})}_{\text{average capacity per route}} \right], \tag{4.12}$$

where $|\mathcal{R}|$ is the number of routes (6), and speed and loading are at their optimum. The left-hand side is the marginal benefit of reallocating capacity to route r, measured as its profit per ton traveled relative to the firm's average. The right-hand side is the marginal cost, measured as the cost of deployment to route r relative to the average.

When $\ell_r = 1$, the firm faces additional marginal cost pressure from the loading constraint. With a fully loaded ship, adding capacity lowers the firm's marginal cost of shipping along

route r (see Footnote 19) relative to its average cost. Since the firm can now fully load a larger ship, its profit on route r rises relative to the other routes.[20]

Entrants in the shipping market. A firm that enters earns the expected profit of active firms, $\left(\int_0^N \pi(j)dj\right)/N$, while staying out delivers zero. As firms enter, congestion rises and shipping prices fall until the route-specific profits are exactly offset by the cost of deploying capacity across routes, leaving the marginal entrant indifferent. The free-entry condition determines the number of shipping firms, N.

Importers. A measure $\mathcal{I}$ of identical importers sources goods from different regions and routes. They buy goods from a shipping firm j operating route r at the price $p_r(j)$ and resell them to consumers at the final price p_o for $o \in \mathcal{O}$. The final price depends only on the origin of the goods, since consumers are indifferent to the route or port by which goods arrive.

Importers sourcing goods from route r produce output with a decreasing-returns-to-scale technology that aggregates all services produced by shipping firms operating along the route, with elasticity of substitution $1/(1-\sigma)$:

$$y_r = \left[\chi_r \left(\int_0^N y_r(j)^\sigma dj\right)^{1/\sigma}\right]^\theta, \tag{4.13}$$

where χ_r is the productivity of importers specializing in route-r goods, and $\theta \in (0,1)$ governs returns to scale from dealing with shipping firms, common across routes.

Importers' profit per route. Importers' profit from each route r is

$$\Pi_r = \max_{\{y_r(j)\}} \left\{ p_o \chi_r^\theta \left(\int_0^N y_r(j)^\sigma dj\right)^{\theta/\sigma} - \int_0^N p_r(j) y_r(j) dj - f \right\}, \tag{4.14}$$

where $f > 0$ is the fixed cost of operating in the retail sector. Since $\theta < 1$, importers can make profits on some routes.

Optimal sourcing decision. We make importers' sourcing decision endogenous, giving shipping disruptions both an intensive and an extensive margin. Disruptions then alter both the quantity importers demand and their distribution across routes. We introduce idiosyncratic

[20] The first-order condition when the optimal load along route r is at the upper bound ($\ell_r = 1$) is

$$\frac{(p_r - mc_r)\, x_r n_r(s_r) \ell_r}{g(K_r)} - \frac{1}{|\mathcal{R}|} \sum_h \frac{(p_h - mc_h)\, x_h n_h(s_h) \ell_h}{g(K_h)}$$
$$= \overline{p}_\kappa \left[\kappa_r - \frac{1}{|\mathcal{R}|}(\kappa - \overline{\kappa})\right] + \frac{\partial mc_r}{\partial \kappa_r} q_r - \frac{1}{|\mathcal{R}|} \sum_h \frac{\partial mc_h}{\partial \kappa_h} q_h,$$

where $\partial mc_h / \partial \kappa_h = -\left(mc_h/\kappa_h\right) \mu_h / \left[\nu \overline{p}_o \kappa_h^\nu + \mu_h\right] < 0$ when $\ell_h = 1$.

taste shocks ϵ_r for goods sourced from route r, assumed mean-zero Type I extreme value with inverse scale parameter $\lambda > 0$. So importers choose the route they source from according to

$$\Pi = \max_{r \in \mathcal{R}} \{\Pi_r + \epsilon_r\}. \tag{4.15}$$

The optimal sourcing decision for goods from route r is governed by the logit expression

$$o_r = \frac{e^{\lambda \Pi_r}}{\sum_{h \in \mathcal{R}} e^{\lambda \Pi_h}}. \tag{4.16}$$

More profitable routes attract more importers.

Importers' demand for shipping services. The demand for goods sourced from route r carried by firm j, $\mathrm{y}_r(j)$, is the per-importer quantity $y_r(j)$ times the mass of importers sourcing route r, $o_r\mathcal{I}$, that is

$$\mathrm{y}_r(j) \equiv y_r(j) o_r \mathcal{I} = \left[\frac{\theta \chi_r^{\sigma} p_o y_r^{(1-\sigma/\theta)}}{p_r(j)} \right]^{1/(1-\sigma)} o_r \mathcal{I}, \tag{4.17}$$

where the final expression replaces $y_r(j)$ with the solution to the importer's problem (4.14).

Entrants in the retail market. Entry occurs before the taste shocks realize. Upon entering, a potential importer draws $\{\epsilon_r\}_{r \in \mathcal{R}}$ and chooses its route according to problem (4.15); staying out yields zero. Free entry drives the expected pecuniary profit of an entrant, $\sum_{r \in \mathcal{R}} o_r \Pi_r = 0$, to zero, determining the measure of importers, $\mathcal{I}$.[21]

Consumers. A unit measure of consumers each holds one unit of time, supplied inelastically at wage w. They also receive a lump-sum transfer Λ from the firms they own. Consumers buy domestic goods d at price p_d and imports c_o from Asia, Europe, the Middle East, and South America ($o \in \mathcal{O}$) at retail price p_o. Thus, the budget constraint of a consumer is:

$$\sum_{o \in \mathcal{O}} p_o c_o + p_d d = w + \Lambda. \tag{4.18}$$

Imported goods from different regions combine into the bundle

$$c = \left(\sum_{o \in \mathcal{O}} \gamma_o c_o^{\rho} \right)^{1/\rho}, \tag{4.19}$$

[21] Under the alternative convention in which entrants also value the taste shocks, the entry condition is the usual logsum. We adopt the pecuniary condition so that the taste shocks affect only the allocation of sourcing but not welfare or households' budget constraint.

where γ_o is an origin-specific utility weight, with $\sum_{o\in\mathcal{O}}\gamma_o = 1$, and $1/(1-\rho)$ is the elasticity of substitution. Goods from the same region sourced via different routes are perfect substitutes.

A consumer derives utility from domestic goods and the imported bundle, with the elasticity of substitution $1/(1-\varrho)$. The consumer solves

$$u = \max_{c_A, c_E, c_M, c_S, d} \left[\gamma_d d^{\varrho} + (1-\gamma_d) c^{\varrho}\right]^{1/\varrho}, \tag{4.20}$$

subject to the budget constraint (4.18) and the imported bundle (4.19).

4.2. *Equilibrium*

Definition. An equilibrium is a solution for: (1) the domestic goods price and domestic labor demand, $\{p_d, l_d\}$; (2) each shipping firm's profit, output, trips, sailing speed, price, load, and capacity along each route $r \in \mathcal{R}$, $\{\pi_r(j), q_r(j), n_r(j), s_r(j), p_r(j), \ell_r(j), \kappa_r(j)\}$, and total profits $\pi(j)$, for all firms $j \in [0, N]$; (3) importers' demand for the services of each shipping firm j, $\mathrm{y}_r(j)$, output y_r, profits Π_r, and sourcing decision o_r for each route $r \in \mathcal{R}$; (4) consumers' demand for domestic goods d and imports c_o from each region $o \in \mathcal{O}$; (5) the wage w; (6) retail prices p_o for each region $o \in \mathcal{O}$; (7) the number of shipping firms N; (8) the number of importers $\mathcal{I}$; (9) aggregate transfers Λ. These satisfy:

1. The domestic firms' problem satisfies (4.1) and their labor demand is $l_d = 1 - \frac{1}{w}\int_0^N \sum_{r\in\mathcal{R}} \left\{\zeta \overline{p}_n n_r(j)^{\zeta} - \frac{\mu_r(j)}{\nu} + \overline{p}_{\kappa}\frac{\kappa_r(j)^2}{2}\right\} dj$.
2. Given the transport price $\overline{p}_n$ and origin goods prices $\overline{p}_o$, each shipping firm j maximizes profits $\pi_r(j)$ per (4.9), yielding $\{q_r(j), n_r(j), s_r(j), p_r(j), \ell_r(j)\}$ along each route r that satisfy (4.3), (4.5), (4.8), (4.10), and $\ell_r(j) = \min\left\{1, \left(\frac{\zeta}{\nu}\frac{\overline{p}_n}{\overline{p}_o}\right)^{1/\nu} n_r(j)^{\zeta/\nu}/\kappa_r(j)\right\}$. Total profits $\pi(j)$ and capacity per route $\kappa_r(j)$ satisfy (4.11) and (4.12).
3. Given shipping prices $p_r(j)$ and retail prices p_o, importers' allocations $\{y_r, \Pi_r, o_r\}$ and $\{\mathrm{y}_r(j)\}_0^N$ satisfy (4.13), (4.14), (4.16), and (4.17) for each route $r \in \mathcal{R}$.
4. Given prices p_d and p_o, the wage w, and profits Λ, consumers' demands for d and c_o, $o \in \mathcal{O}$, maximize utility (4.20) subject to the budget constraint (4.18) and bundle (4.19).
5. The shipping-importer market clears on each route r according to $\int_0^N q_r(j)dj = o_r\mathcal{I}\left[\theta\chi_r^{\sigma} p_o y_r^{(1-\sigma/\theta)}\right]^{1/(1-\sigma)} \int_0^N p_r(j)^{1/(\sigma-1)}dj$.
6. The importer-consumer market clears: $c_A = (o_{AEP}y_{AEP} + o_{AES}y_{AES} + o_{AW}y_{AW})\mathcal{I}$ for Asian goods, and $c_E = o_{EE}y_{EE}\mathcal{I}$, $c_M = o_{ME}y_{ME}\mathcal{I}$, and $c_S = o_{SE}y_{SE}\mathcal{I}$ for the single-route regions, yielding retail prices p_o.
7. Free entry in shipping holds, $\int_0^N \pi(j)dj = 0$, determining N.
8. Free entry in retail holds, $\sum_{r\in\mathcal{R}} o_r\Pi_r = 0$, determining $\mathcal{I}$.
9. Retail entry costs are rebated to consumers, $f\mathcal{I} = \Lambda$.

4.3. *Unpacking the impact of shipping disruptions*

Akin to Hulten (1978)'s theorem (see Baqaee and Farhi, 2019, 2020), Proposition 4.1 provides closed-form expressions for how shipping disruptions along a route transmit to aggregate output and consumer welfare. The impact depends on the route's importance in aggregate output (its Domar weight), with amplifying or dampening effects through the retail sector via importers' output and sourcing decisions. Because a disruption alters the profits that draw importers into the retail market, there is another channel that operates through entry and exit.

Let $X_r \equiv \left[\int_0^N x_r(j)^{\sigma/(1-\sigma)} dj\right]^{(1-\sigma)/\sigma}$ denote aggregate route-r shipping capacity utilization net of congestion, bundling the exogenous primitives $x_r(j)$ across the endogenous measure of shipping firms N.[22] A fall in X_r, driven by a deterioration in $x_r(j)$, is an adverse shipping disruption along route r. Our measured route-r utilization that includes the endogenous effect of congestion is $Z_r \equiv X_r / g(K_r)$, which coincides with $\left[\int_0^N z_r(j)^{\sigma/(1-\sigma)} dj\right]^{(1-\sigma)/\sigma}$. The proposition below can be stated equivalently in terms of X_r or Z_r.

PROPOSITION 4.1: (IMPACT OF SHIPPING DISRUPTIONS ON REAL AGGREGATE OUTPUT) The elasticity of real aggregate output Y with respect to route-r utilization, X_r, is

$$\frac{\partial Y}{\partial X_r}\frac{X_r}{Y} = \sum_{h\in\mathcal{R}} \underbrace{\left(\begin{smallmatrix}\text{Route } h\text{'s share in}\\ \text{aggregate output}\end{smallmatrix}\right)}_{\text{Domar weight}} \times \underbrace{\left(\underbrace{\frac{\partial y_h}{\partial X_r}\frac{X_r}{y_h}}_{\text{intensive margin}} + \underbrace{\frac{\partial o_h}{\partial X_r}\frac{X_r}{o_h}}_{\text{sourcing reallocation}} + \underbrace{\frac{\partial \mathcal{I}}{\partial X_r}\frac{X_r}{\mathcal{I}}}_{\text{retail entry}} \right)}_{\textit{endogenous}\text{ trade elasticity}} + \left(\begin{smallmatrix}\text{Domestic share in}\\ \text{aggregate output}\end{smallmatrix}\right) \times \underbrace{\left(\frac{\partial d}{\partial X_r}\frac{X_r}{d}\right)}_{\text{domestic spillovers}},$$

where route h's share in aggregate output is $p_{o(h)} y_h o_h \mathcal{I} / (w + \Lambda)$, with $o(h)$ the origin region of route h and $p_{o(h)}$ the retail price of goods from that origin. An analogous expression holds for the impact of the congestion externality $g(K_r)$ or measured utilization Z_r on output.

PROOF: See Appendix B.1. *Q.E.D.*

The proposition has a few important implications. Hulten's (1978) theorem reduces the first-order effect of a route-r productivity shock on output to route r's Domar weight alone, since the envelope theorem kills every reallocation term once marginal value products are equalized. Our decomposition nests that benchmark. For example, the equilibrium is efficient when we have exogenous iceberg costs (i.e., no congestion externality, no shared capacity to reallocate, and shipping priced at its social marginal cost). Cross-route substitution, sourcing, entry, and the shift between domestic and imported goods then do not matter to first order, and the elasticity of aggregate output collapses to route r's Domar weight. By contrast, in our benchmark specification, each distortion generates first-order welfare effects.

[22] Although the $x_r(j)$ are exogenous, X_r is endogenous because N adjusts in equilibrium.

First, a shock to route r enters the term of *every* route h in the sum, not only its own, through two channels. On the demand side, a disruption raises the price of goods arriving along route r, pushing consumers and importers toward goods from other routes or regions. On the supply side, shipping firms reallocate their capacity across the six existing routes. Hence, a disruption in one route changes congestion and marginal costs in all the others.

Second, the equilibrium is inefficient. Imported goods sell above marginal cost—reflecting the shipping markup $(1/\sigma)$ and the importers' decreasing-returns margin $(1/\theta)$—so reallocation between imports and domestic goods changes welfare at the margin. And because no shipping firm internalizes the congestion externality, $g(K_r)$, the allocation of capacity across routes is distorted as well. Absent the externality, the cross-route reallocation terms would cancel to first order as the retail markup is identical across routes.

Third, importers convert shipped goods into retail output with diminishing returns. As a result, they absorb a disruption only partially. The residual adjustment operates through sourcing reallocation and the entry and exit of importers. Together, these forces imply that route r's Domar weight is *not* a sufficient statistic for the welfare impact of a disruption.

These channels also distinguish our benchmark from canonical quantitative trade models. Arkolakis et al. (2012) show that, across a class of models, the welfare effect of a change in trade frictions is captured by two sufficient statistics, the domestic expenditure share and the trade elasticity. The result relies on iceberg trade costs and an efficient economy. Subsequent work shows that this sufficiency can break in richer environments (e.g., with firm heterogeneity, production networks, Melitz and Redding, 2015, Baqaee and Farhi, 2019, 2020). Our framework departs on both counts. First, trade costs are not exogenous. They are the equilibrium price of a shipping service whose capacity is reallocated across a shared, congestible network, so a route shock has spillovers that a single trade elasticity cannot capture. Second, the equilibrium is inefficient, the joint result of congestion externalities, markups, and the sourcing externality. The domestic share is not a sufficient statistic, since the same change in the import share can map to different welfare changes depending on which margin adjusts.

Propositions B.1 and B.2 in Appendix B show how shipping disruptions affect importers' output and sourcing decisions in partial equilibrium. Section 6 uses the model to compute the route-specific Domar weights and the elasticity of importers' output and sourcing to a shipping productivity shock along each route, taking their general equilibrium effects into account.

4.4. *The constrained social planner's problem*

The decentralized equilibrium has three sources of inefficiency. First, individual shipping firms do not internalize the congestion externality $g(K_r)$ when they allocate capacity. Second, shipping firms have monopoly power and charge a markup $m_r = 1/\sigma$ over marginal cost. Third,

decreasing returns in importers' technology ($\theta < 1$) together with idiosyncratic sourcing tastes generate retail profits that determine entry and sourcing incentives across routes, and shipping firms ignore those incentives when setting prices.

We now study the efficient allocation chosen by a constrained social planner and compare it with the decentralized equilibrium. The planner chooses imported goods from each region o, $\{c_o\}_o$, labor in domestic production, l_d, and, for each route r, the aggregate shipping capacity, loading, average speed, social price of shipping, and importers' demand, $\{K_r, \ell_r, s_r, v_r^{\text{social}}, Y_r^{\text{input}}\}_r$, to maximize consumer welfare

$$\max_{\substack{\{c_o\}_{o\in\mathcal{O}}, l_d, \\ \{K_r,\ell_r,s_r,v_r^{\text{social}},Y_r^{\text{input}}\}_{r\in\mathcal{R}}}} \mathbb{E}_{\boldsymbol{\epsilon}}[u] \quad \text{with } d = z_d l_d. \tag{4.21}$$

The planner respects the aggregate capacity constraint, $\sum_{r\in\mathcal{R}} K_r + \overline{K} = K$, the free-entry conditions in retail and shipping that determine the masses of importers and shipping firms, $\mathcal{I}$ and N, and the distribution of importers' taste shocks across routes. As in the decentralized equilibrium, the planner takes as given the disruption and taste shocks, domestic productivity, and the transportation, capacity-adjustment, and foreign-goods prices. The planner internalizes the congestion externality and how importers reallocate across routes. Appendix B.3 states the problem in full.

The planner allocates capacity across routes in two different ways. First, the first-order condition yields an expression analogous to Equation (4.12), summed over all N shipping firms. Because the planner internalizes the congestion externality, she scales the marginal cost of shipping along route r up by the factor $(1+\phi)$, which lowers the benefit of reallocating ships across routes. Second, to cover the cost of redeploying ships across routes, the planner sets a social price of shipping on each route that includes a markup. But this markup does not depend on the elasticity of substitution between shipping services, σ. The relevant elasticity is now θ, the elasticity of retail output with respect to total import demand. And because the planner also internalizes how shipping prices affect importers' sourcing decisions, the markup includes a further term for the sourcing response.

5. TAKING THE MODEL TO DATA

We match the model to U.S. data to capture the shipping trends in Sections 2 and 3. We calibrate the model to a pre-disruption economy using February 2016 data, then simulate from March 2016 through March 2025 (weekly data are aggregated to a monthly frequency).

Parameters. The framework has 42 parameters: 7 consumer preference parameters, $\{\rho, \varrho, \gamma_d, \gamma_o\}$ for $o \in \mathcal{O}$; 10 governing the importers' problem, $\{\chi_r, \theta, \sigma, f, \lambda\}$ for $r \in \mathcal{R}$; 18

for the shipping market, $\{x_r, A_r, \varepsilon_k, \phi, \alpha, \zeta, \nu, \overline{\kappa}\}$ for $r \in \mathcal{R}$; the productivity of domestic producers, z_d; and 6 exogenous prices (capacity redeployment, transportation, and foreign goods), $\{\overline{p}_\kappa, \overline{p}_n, \overline{p}_o\}$ for $o \in \mathcal{O}$. We split these parameters into two groups. The first have direct empirical counterparts, are estimated outside the model, or can be normalized. The second we back out from the equilibrium conditions to match specific data targets exactly. Table 5.1 summarizes the full set of estimated parameters. Table 5.2 compares the targeted model moments with the data as of February 2016.

5.1. *Exogenous and externally estimated parameters*

Consumer elasticity of substitution. A large literature estimates trade elasticities of substitution. A robust finding is that the elasticity of substitution across goods from different countries exceeds the elasticity between domestic and foreign goods: consumers substitute more readily among import sources than away from domestic varieties. Feenstra et al. (2018), using U.S. data from 1992 to 2007, find a "micro-Armington" elasticity (across foreign goods) of 3.22 to 4.05 and a "macro-Armington" elasticity (domestic versus foreign) near 2 for most sectors. We set the substitution parameter between imported goods to $\rho = 0.75$ (elasticity of 4) and between foreign and domestic goods to $\varrho = 0.5$ (elasticity of 2).

Shipping technology and costs. We estimate α, the elasticity of the number of trips to sailing speed, directly from the data. Satellite-based AIS data from January 2016 through March 2025 show that intercontinental voyages often take up much of a calendar month, leaving observed monthly trip counts clustered at small integer values. To recover α without censoring bias, we reconstruct the average number of trips per ship from the DWT-weighted average trip duration (in hours) along route r in month t, $\overline{\texttt{Number of Trips}}_{r,t} = 730/\overline{\texttt{Trip Duration}}_{r,t}$, where 730 is the approximate number of hours in a month and the average trip duration is $\overline{\texttt{Trip Duration}}_{r,t} = \sum_{i \in \boldsymbol{\nu}_t} \left(\frac{\mathbf{1}_{i,r,t} \times \texttt{DWT}_i}{\sum_{i' \in \boldsymbol{\nu}_t} \mathbf{1}_{i',r,t} \times \texttt{DWT}_{i'}} \right) \times \texttt{Trip Duration}_{i,r,t}$, and $\texttt{Trip Duration}_{i,r,t}$ is the average number of hours vessel i spends on a trip on route r in month t. The monthly DWT-weighted average sailing speed follows from Equation (2.2).

To recover α, we regress the average number of trips on average sailing speed in logs. Table A1 in Appendix C.1 reports the full results, comparing OLS with WLS specifications that use total route DWT as analytical weights to dampen idiosyncratic monthly noise. We set $\alpha = 0.663$, the estimate from the WLS specification with route, year, and month fixed effects.

Shipping costs are increasing and strictly convex in both the number of trips and the quantity transported per trip. Lacking granular firm-level data to estimate these elasticities, we impose a quadratic cost structure and set the trip elasticity to $\zeta = 2$. The constant-returns restriction, $1/\zeta + 1/\nu = 1$, implies that the elasticity to the transported quantity is $\nu = 2$.

Congestion. We estimate the congestion exponent ϕ as the elasticity of total congestion time with respect to the total number of port calls, which coincides with the elasticity of congestion with respect to deployed capacity when trips per vessel do not vary systematically with the capacity serving the port. Using vessel-level AIS data and the IMA-DBSCAN algorithm from Section 2, we build a monthly panel of total port calls and total congestion times for major global container ports from January 2016 through March 2025. We compute the total hours container ships wait in a port's anchorage before berthing, then regress this on the total number of port calls in logs. Table A2 in Appendix C.1 reports the results. The preferred WLS specification, weighting by total port calls and controlling for port, year, and month fixed effects, yields a precisely estimated elasticity of 1.296. The estimate is stable under year-by-month fixed effects and rises when 2020–2022 is excluded (Table A2), so we view it as conservative; the welfare estimates for the historical episodes are in any case invariant to ϕ (Appendix C.7). We set $\phi = 1.296$ and normalize the congestion shifter ε_k to 1.

Importer technology and sourcing. We normalize importer productivities to one except on two of the three Asian routes. Because goods from the same origin are perfect substitutes, the relative productivities of the Asian routes must be calibrated internally to match observed shipping capacity shares. We jointly estimate χ_{AES} and χ_{AW} (more below) and normalize $\chi_{AEP} = \chi_{EE} = \chi_{ME} = \chi_{SE} = 1$.

The price elasticity of importers' demand for shipping services is $\xi_r = 1/(1-\sigma)$, implying a shipping-firm markup of $1/\sigma$. We target a 25% shipping markup, which pins down $\sigma = 0.8$. The marginal cost of importers conditional on sourcing a route is θp_o, yielding a markup of $1/\theta$. We set $\theta = 0.8$ to match a 25% retail markup in line with the evidence in Marto (2025). Finally, we normalize the inverse scale parameter of the mean-zero Gumbel sourcing shock to $\lambda = 1$.[23] Lower values of λ raise the variance of importers' idiosyncratic taste shocks and give a more dispersed sourcing distribution.

Domestic producers. We normalize the productivity of domestic producers to $z_d = 1$, so the price of domestic goods equals the wage. With the wage as numeraire, $p_d = 1$.

5.2. *Internally calibrated parameters*

Import shares. The consumption weights on domestic and imported goods, $\{\gamma_d, \gamma_A, \gamma_E, \gamma_M, \gamma_S\}$, help match the import share from each region as of February 2016. Regional imports are taken from the Census Bureau as the containerized vessel values and charges from a region.[24] In

[23] In principle, λ could be identified from the elasticity of sourcing shares to route profitability, but this requires importer-level data linking products to routes that are not yet available to us. We leave this for future work, but we assess the sensitivity of our results to this choice.

[24] The Census Bureau trade data are available here and GDP is taken from the BEA (here). Countries are assigned to regions following the World Bank's classification (here).

TABLE 5.1

PARAMETER VALUES

Parameter	Description	Value	Identification
Consumer preferences			
ϱ	Elast. of sub. between imported and domestic goods	0.500	Feenstra et al. (2018)
ρ	Elast. of sub. among imported goods	0.750	Feenstra et al. (2018)
γ_d	Cons. weight, domestic goods	0.985	Total import share
γ_o	Cons. weight, Asia, Europe, Middle East, South America	0.369, 0.356 0.093, 0.182	Regional import shares
Retail market			
χ_r	Importer prod., Asia to East Coast via Panama/Suez/to West Coast	1.000, 1.084, 1.238	Potential shipping capacity,
	Europe, Middle East, South America to East Coast	1.000, 1.000, 1.000	normalization
θ	Returns to scale	0.800	Retail markup
σ	Elast. of sub. between shipped goods	0.800	Shipping markup
f	Entry cost in retail market	0.008	Number of retail firms
λ	Inverse scale, sourcing taste shocks	1.000	Normalization
Shipping market			
x_r	Route prod., Asia to East Coast via Panama/Suez/to West Coast	2.760, 3.493, 5.642	Shipping disruptions
	Europe, Middle East, South America to East Coast	3.240, 1.558, 2.110	
A_r	Speed eff., Asia to East Coast via Panama/Suez/to West Coast	0.300, 0.383, 0.564	Average speed
	Europe, Middle East, South America to East Coast	0.403, 0.107, 0.263	
ε_k	Congestion shifter	1.000	Normalization
ϕ	Congestion exponent	1.296	Congestion time
α	Exponent on speed	0.663	Elast. trip number wrt speed
ζ	Exponent on number of trips	2.000	Exogenous
ν	Exponent on quantity transported	2.000	Constraint
$\overline{\kappa}$	Idle capacity	0.721	Non-US-bound capacity
$\overline{p}_\kappa$	Redeployment price	0.009	Number of shipping firms
$\overline{p}_n$	Transportation price ($\times 100$)	0.001	Average number of trips
Producers			
z_d	Productivity of domestic firms	1.000	Normalization
$\overline{p}_o$	Origin prices, Asia, Europe, Middle East, South America	0.019, 0.041 0.013, 0.016	Potential shipping capacity

February 2016, the containerized import share of GDP was 3.9%, with 2.9% from Asia, 0.7% from Europe, 0.001% from the Middle East, and 0.3% from South America.

Given equilibrium prices and the consumer elasticities, we back out γ_d and $\{\gamma_o\}_{o\in\mathcal{O}}$ to match these four targets exactly, subject to $\sum_{o\in\mathcal{O}}\gamma_o = 1$. Solving the consumer's problem in (4.20) gives the expenditure share on domestic goods, $\omega_d \equiv p_d d/(w+\Lambda) = \gamma_d^{1/(1-\varrho)}\left(\frac{p_d}{P^{\text{total}}}\right)^{\varrho/(\varrho-1)}$, and the share on imports from region $o\in\mathcal{O}$, $\omega_o \equiv p_o c_o/(w+\Lambda) = \gamma_o^{1/(1-\rho)}\left(\frac{p_o}{P^{\text{imp}}}\right)^{\rho/(\rho-1)}(1-\omega_d)$. Here $P^{\text{total}} \equiv \left[\gamma_d^{1/(1-\varrho)}p_d^{\varrho/(\varrho-1)} + (1-\gamma_d)^{1/(1-\varrho)}\left(P^{\text{imp}}\right)^{\varrho/(\varrho-1)}\right]^{(\varrho-1)/\varrho}$ is the economy-wide ideal price index and $P^{\text{imp}} \equiv \left(\sum_{o\in\mathcal{O}}\gamma_o^{1/(1-\rho)}p_o^{\rho/(\rho-1)}\right)^{(\rho-1)/\rho}$ the price index for imports.

Shipping disruptions. Route-specific shipping utilization, z_r, is from Section 3 and Figure 3.5b. Given the endogenous congestion externality $g(K_r)$, we back out the exogenous component of shipping utilization, $\{x_r\}_{r\in\mathcal{R}}$, to match the empirical ratio exactly.

Potential shipping capacity. The origin prices $\{\overline{p}_A, \overline{p}_E, \overline{p}_M, \overline{p}_S\}$ and the retail productivities for Asia to the East Coast via Suez and to the West Coast, χ_{AES} and χ_{AW}, help match potential shipping capacity (in DWT) by route as of February 2016 (see Figure 2.3b). The theoretical counterpart is $\int_0^N n_r(j)\ell_r(j)\kappa_r(j)dj$. Given the speed and cost exponents (α,ν,ζ), speed ef-

ficiencies (A_r), sailing speeds (s_r), and the mass of shipping firms (N), the origin prices and retail productivities exactly match the observed route-specific tonnages.

Average sailing speeds. We identify the speed efficiency shifters $\{A_r\}_{r\in\mathcal{R}}$ from the average sailing speed on each route (Equation (2.2)) in February 2016 (Figure 3.4a). Given the targeted speed s_r and the optimal number of trips n_r per route, A_r is uniquely identified.

Average number of trips. The price per trip, $\overline{p}_n$, makes the model's average number of trips per firm per route per month match the AIS data. The theoretical average is $\overline{n} = \sum_{r\in\mathcal{R}} \int_0^N n_r(j)dj / [N|\mathcal{R}|]$. In February 2016, a vessel averaged 2.021 U.S.-bound trips per route per month, with trip lengths ranging from 7 to 26 days.

Idle capacity. Idle capacity, $\overline{\kappa}$, matches the fraction of global capacity allocated to U.S.-bound routes, that is $\varpi_{US} = \int_0^N \sum_{r\in\mathcal{R}} \kappa_r(j)dj / \left[\int_0^N \sum_{r\in\mathcal{R}} \kappa_r(j)dj + N\overline{\kappa}\right]$. We use AIS data to track the global fleet and define the target as the ratio of U.S.-bound potential capacity (inbound and outbound) to total global potential capacity. In February 2016, 59.8% of the world's DWT was allocated to non-U.S.-bound routes and the rest to U.S.-bound routes.

Number of importers and shipping firms. The retail entry cost, f, normalizes the measure of importers, $\mathcal{I}$, to one. The capacity-redeployment price, $\overline{p}_\kappa$, helps match the average number of shipping firms, N, across U.S.-bound routes. In February 2016, the average was 33 firms.

TABLE 5.2

TARGETED MOMENTS: DATA VS. MODEL (FEBRUARY 2016)

Moment	Model	Data	Source
Import shares of output, Total	0.039	0.039	Census
Asia, Europe, Middle East, South America	0.029, 0.007, 0.000, 0.003	0.029, 0.007, 0.000, 0.003	
Shipping utilization			
Asia to East Coast via Panama/Suez/to West Coast	0.849, 0.882, 0.953	0.849, 0.882, 0.953	AIS
Europe, Middle East, South America to East Coast	0.823, 0.821, 0.797	0.823, 0.821, 0.797	
Potential shipping capacity (in DWT), Total	24.143	24.143	AIS
Asia to East Coast via Panama/Suez/to West Coast	2.921, 4.714, 11.056	2.921, 4.714, 11.056	
Europe, Middle East, South America to East Coast	3.202, 0.478, 1.773	3.202, 0.478, 1.773	
Sailing speeds			
Asia to East Coast via Panama/Suez/to West Coast	15.433, 15.325, 16.264	15.433, 15.325, 16.264	AIS
Europe, Middle East, South America to East Coast	14.018, 15.824, 12.149	14.018, 15.824, 12.149	
Average number of trips	2.021	2.021	AIS
Non-U.S.-bound capacity (% of world capacity)	0.598	0.598	AIS
Number of shipping firms	33.0	33.0	AIS
Number of importers	1.0	1.0	Normalization

5.3. *Tracing the evolution of U.S.-bound shipping between 2016 and 2025*

We simulate the model at a monthly frequency from March 2016 to March 2025. At each date, we jointly solve for the time-varying exogenous processes that minimize the distance between model-implied aggregates and the target data. In particular, the exogenous components of utilization $x_{r,t}$ match route-level shipping utilization; the speed efficiencies $A_{r,t}$ match average

sailing speeds for each route; origin prices $\overline{p}_{o,t}$ match import shares of output by region; the retail productivities $\chi_{r,t}$ match route-level potential capacity in DWT; idle capacity $\overline{\kappa}_t$ matches the world shipping capacity in DWT; the redeployment price $\overline{p}_{\kappa,t}$ matches the number of shipping firms; and the retail entry cost f_t ensures the number of importers is equal to one. All other parameters are held constant over time. These simulations are the baseline for the counterfactual experiments in Sections 6 and 7.

Figure A1 in Appendix C.1 presents the targeted time series, Figure A2 the exogenous processes that match the data, and Figure A3 other aggregates of interest.

The simulation attributes the 2020–2022 disruption period to a combination of forces. The congestion externality, $g(K_r)$, rises markedly on most routes from summer 2020 to the end of 2022, accounting for part of the decline in shipping utilization, z_r; the remainder is assigned to the exogenous component, x_r (Panels A1a, A2a, and A3a). The rise in congestion reflects the surge in U.S.-bound capacity as idle capacity, $\overline{\kappa}$, falls and vessels redeploy toward U.S. routes. Together with a lower capacity-redeployment price, these rationalize the entry of shipping firms and the rise in world capacity in 2021–2022 (Panels A1b, A1c, A1d, A2b, A2c, and A3b).

The simulation translates the 2021–2022 disruptions into about four additional days at sea along the Asian routes and higher loading factors for vessels departing Asia relative to other routes, while trips per ship change little (Panels A3c, A3d, and A3e). The observed decline in sailing speeds in 2021 is driven by higher speed efficiencies, A_r (Panels A1e and A2d).

Matching the decline in the Asian import share during the COVID-19 period requires retail productivities, χ_r, to fall during the pandemic and recover quickly afterward (Panels A1f and A2f). In the model, the lower Asian share is accompanied by higher retail prices for Asian goods, smaller Asian route shares of output, and a reallocation of importers' sourcing toward other routes (Panels A3g, A3h, and A3i).

Validation. The simulation targets quantities and utilization but no prices, so freight rates provide a check on the model's price mechanism that the calibration cannot deliver by construction. In the model, the marginal cost of shipping is inversely proportional to utilization (Equation 4.8), so shipping prices should rise one-for-one with shipping pressures, measured as $-\log z_{r,t}$. Two recent disruption episodes provide descriptive evidence consistent with this implication. During the pandemic, the SCFI (Shanghai Containerized Freight Index) spot rate from Shanghai to the U.S. West Coast rose from \$1,572 per FEU in January 2020 to \$8,079 in February 2022, as our measured shipping capacity utilization on the Asia–West Coast route fell from over 90% before the pandemic to 53% in late 2021. During the Red Sea crisis, the Shanghai–U.S. East Coast rate rose from \$2,847 per FEU in December 2023 to \$9,784 in July 2024, as rerouting reduced our measured shipping capacity utilization on the Asia-East Coast route via the Suez Canal to 61% (Clarksons Research Services Limited, 2025).

Appendix C.2 presents regressions of log spot rates on shipping pressure at the route-month level. The elasticities are positive and statistically significant in both episodes, even though the episodes differ in origin (port congestion driven by both demand and supply pressures in one case and a security shock in the other). Spot rates reflect only the freight component, whereas $p_{r,t}$ is a delivered price inclusive of the goods, so we interpret the comovement as validating the sign and timing of the model's price response rather than its level.

6. The macroeconomic impact of shipping disruptions

Using the calibrated model, we revisit Proposition 4.1 and quantify the macroeconomic effects of route-specific shipping disruptions. We then focus on two disruption events. The first is the unprecedented global port congestion of summer 2021, which disproportionately affected the Asia–West Coast route. The second is the rerouting away from the Suez Canal following the 2023–2024 Houthi attacks in the Red Sea, which significantly affected the routes from Asia via the Suez Canal and from the Middle East. In both scenarios, counterfactual economies without the measured disruptions are simulated and compared to our baseline economy.

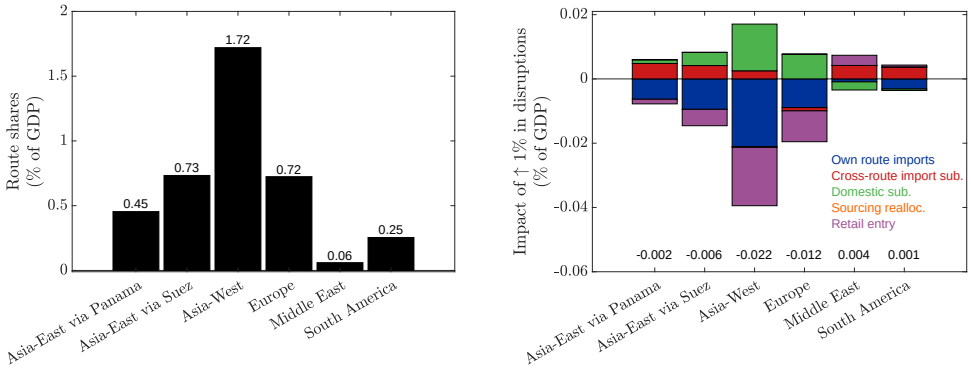
6.1. *The aggregate impact of route-specific disruptions*

We use the calibrated model to quantify the decomposition in Proposition 4.1, computing the elasticity of aggregate output with respect to a route-specific disruption shock (that is, a fall in the exogenous component of route-r utilization, X_r) accounting for general equilibrium effects. We evaluate the elasticities as of February 2016, when routes operated near their historical utilization levels, so that they capture the transmission mechanism in an undisrupted network rather than in a stressed one. These elasticities are state-dependent and their magnitudes and signs can vary across dates, particularly during disruption episodes.

Figure 6.1a shows each route's share in U.S. output (its Domar weight) as of February 2016. The Asia–West Coast route accounted for nearly 1.7% of GDP, while the Asia–East Coast route via the Suez Canal and the Europe–East Coast route each accounted for about 0.7%. The route weights sum to the containerized import share (3.9% of GDP).

Figure 6.1b reports the elasticity of aggregate output to shipping disruptions along each of the six routes. A 1% fall in route r's utilization, X_r, reduces aggregate output by about 0.022% for a disruption along the Asia–West Coast route, 0.012% along the Europe–East Coast route, and 0.006% along the Asia–East Coast route via the Suez Canal. The Asia–West Coast elasticity

FIGURE 6.1.—The impact of shipping disruptions on aggregate output



(a) Route shares in aggregate output

(b) Elasticity of output to shipping disruptions

Note: Panel (a) shows each route's share of GDP (its Domar weight). Panel (b) shows the elasticity of output to a route-r disruption shock, $-\frac{\partial Y}{\partial X_r}\frac{X_r}{Y}$. The elasticities account for general equilibrium effects and are computed with automatic differentiation.

is roughly 1.3 times the route's Domar weight: the general equilibrium reallocation amplifies, rather than dampens, the direct effect of a disruption.[25]

For the Middle East and South America routes, the effect runs in the opposite direction: a disruption on these routes slightly raises aggregate output (by 0.004% and 0.001%, respectively, for a 1% fall in utilization). Both routes carry small Domar weights and a disruption there releases capacity that shipping firms redeploy toward the more profitable Asian routes (at the low congestion levels of February 2016, the increased congestion does not overturn this gain).[26]

Figure 6.1b also decomposes each elasticity into the margins of Proposition 4.1: the change in imports from the disrupted route, the substitution toward domestic goods and imports from other routes, importers' sourcing decisions, and retail entry. Disruptions reduce imports on the affected route as importers shift to goods from other routes and consumers to domestic goods, and they further reduce output through importer exit. For the Asia–West Coast route, the fall in the route's imports and the retail exit reduce output by 0.021% and 0.018%, respectively, while substitution toward other goods offsets the impact on output by 0.017%.

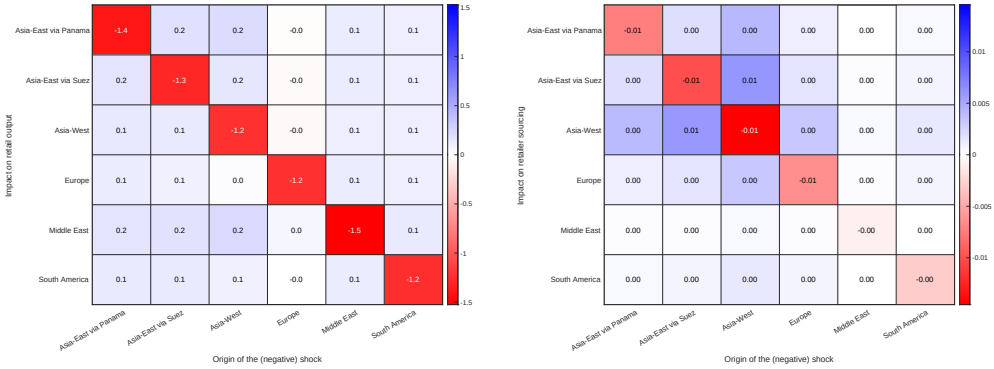
We now zoom in on imports. Figure 6.2a shows the impact of a route-r disruption shock on retail output, $-(\partial \log y_h / \partial \log X_r)$, and Figure 6.2b displays the impact on importers' sourcing decisions, $-(\partial \log o_h / \partial \log X_r)$. A 1% disruption shock reduces retail output from the disrupted route by about 1.2 to 1.5% (diagonal cells) and raises output on competing routes by up

[25]The elasticities fully capture general-equilibrium effects. We compute the partial derivatives of the equilibrium conditions using automatic differentiation and recover the full general-equilibrium Jacobian via the implicit function theorem. This approach yields accurate numerical derivatives.

[26]A disruption acting as a second-best improvement is the same logic that reappears, at a policy scale, in the tariff experiments of Section 8.

to 0.2% (off-diagonal cells), with shocks to the Asian routes generating the largest spillovers. The own-route elasticities are well below the partial equilibrium benchmark of $\theta/(1-\theta) = 4$. The general equilibrium response of retail prices and congestion absorbs roughly two-thirds of the direct effect, in line with Proposition B.1.[27] A 1% shock reduces the fraction of importers sourcing the affected route by up to 0.01% and reallocates them toward competing routes.

FIGURE 6.2.—The effect of shipping disruptions on importers' output and sourcing decisions



(a) Importers' retail output (b) Importers' sourcing decisions

Note: Panel (a) displays the elasticity of importers' route-h output to a route-r disruption shock, $-(\partial \log y_h / \partial \log X_r)$. Panel (b) displays the elasticity of importers' route-h sourcing to a route-r disruption shock, $-(\partial \log o_h / \partial \log X_r)$.

6.2. *Port congestion in the summer and fall of 2021*

The summer and fall of 2021 saw a record number of ships waiting at the ports of Los Angeles and Long Beach, putting massive pressure on the shipping sector, in particular along the Asia–West Coast route. Although port congestion only became news in the summer of 2021, our satellite-based measure signaled a fall in shipping utilization as early as November 2020, and by November 2021 utilization on the Asia–West Coast route had fallen 20 p.p. relative to the prior year, reaching its trough at the end of 2021.

To study the aggregate impact of the 2021 shipping disruptions along the Asia–West Coast route, we build a counterfactual economy in which the route's shipping utilization, z_{AW}, is held at its January 2020 value of 0.880 (about 1.5 times its end-2021 trough) in every month from January to December 2021 (the window over which we evaluate welfare). In each month, we solve for the exogenous component x_{AW} that delivers this value, keeping all other parameters at

[27] Figures A4a and A4b in Appendix C.3 show that a 1% disruption shock raises the retail price of goods from the affected region by 0.05 to 0.52%, with shocks to the Asian routes spilling over into higher prices for all regions. The shock also lowers congestion on the affected route while raising it on the competing routes.

their baseline levels.[28] As shipping firms reallocate capacity onto the route in the counterfactual, x_{AW} must rise enough to offset the induced congestion. The welfare cost we report is therefore the cost of the utilization collapse as a whole, and it should be read as an upper bound on the cost attributable to the exogenous component alone. Figure A5 in Appendix C.4 shows the evolution of the key aggregates, and Table 6.1 summarizes the results.

Without the disruptions, shipping firms find it optimal to reallocate ships from other Asian routes to the Asia–West Coast route, while capacity allocated to Europe, the Middle East, and South America remains largely unchanged, leaving more ships idle. Firms increase the number of voyages along the route from about 3 to 4 trips per ship per month, at the expense of fewer trips along the other Asian routes. U.S.-bound potential shipping supply rises from 47 to 62 million metric tons at the end of 2021 and, because utilization is higher, effective shipping supply rises by more, from 33 to 53 million metric tons. The additional supply is accompanied by a 31% fall in shipping service prices on the Asia–West Coast route.

With more shipping supply, imports rise. The import share of GDP at the end of 2021 increases from 4.1% to 5.2%, driven by Asian goods: their retail price falls by 23%, sourcing the Asia–West Coast route becomes more profitable, and more importers specialize in it. Higher retail profits draw new importers into the market and the measure of importers expands by 20% at end-2021. The aggregate price index falls by much less (1.1% in December 2021 relative to the baseline) as most consumption is domestically produced.

TABLE 6.1

THE MACROECONOMIC IMPACT OF THE 2021 AND 2024 SHIPPING DISRUPTIONS

	2021 port disruptions (December 2021)		2024 Red Sea crisis (May 2024)	
	Baseline	Counterfactual	Baseline	Counterfactual
Panel A: Quantities				
Shipping utilization, affected route (%)[a]	60.5	88.0	59.4	85.7
Potential shipping capacity, total (million metric tons)	47.1	61.8	42.1	47.2
Effective shipping capacity, total (million metric tons)	33.2	53.2	32.5	39.0
Trips per ship per month, affected route[a]	3.1	3.8	1.6	2.7
Import share, total (% of GDP)	4.1	5.2	3.7	4.0
Measure of importers	1.0	1.2	1.0	1.1
Panel B: Prices (% deviation of counterfactual from baseline)				
Shipping price, affected route[a]	−31.2		−30.7	
Retail price, Asian goods	−23.3		−9.6	
Aggregate price index	−1.1		−0.3	
Welfare: equivalent variation (% of output)[b]	0.69		0.35	

Notes: The table compares the baseline economy with the counterfactual economies, evaluated in December 2021 (West Coast port disruptions) and May 2024 (Red Sea crisis). [a] For the 2024 episode, the affected route shown is Asia–East Coast via the Suez Canal. On the Middle East–East Coast route, utilization is 50.1% in the baseline and 80.9% in the counterfactual, trips rise from 0.6 to 0.8 per ship per month, and the shipping price falls by 38.1%. [b] The equivalent variation is computed as in Equation (6.1) over January–December 2021 for the West Coast port disruption and over December 2023–March 2025 for the Red Sea crisis.

[28] As each month is a static equilibrium, months outside the window are unaffected by the experiment.

The welfare impact of the 2021 West Coast port disruptions is the value of ε_{ev} that solves

$$(1+\varepsilon_{\text{ev}})\sum_{\tau=t}^{T}\beta^{\tau-t}u_{\tau}^{\text{Base}}=\sum_{\tau=t}^{T}\beta^{\tau-t}u_{\tau}^{\text{Exp}}, \tag{6.1}$$

where the left-hand side is utility in the baseline economy, the right-hand side utility in the counterfactual economy, $\beta = 0.9966$ is the consumer's monthly discount factor (equivalent to an annualized discount rate of 4%), and the window runs from $t =$ January 2021 to $T =$ December 2021. The equivalent variation ε_{ev} is the change in consumer income, at baseline prices, that would make the consumer as well off as in the counterfactual economy without shipping disruptions. Since utility is higher without disruptions, consumers would pay a positive amount to live in that economy: about 0.69% of output to avoid the 2021 West Coast port disruptions.

Figure 6.3a decomposes the welfare impact into the margins of Proposition 4.1.[29] The 2021 West Coast port disruptions operate mainly through own-route imports, reinforced by retail entry and partially offset by the substitution away from domestic goods and other imported goods.

Figure 6.3b plots year-over-year CPI inflation alongside the counterfactual inflation rate implied by the model's aggregate price index deviation, i.e., the inflation rate that would have prevailed had the disruption not occurred, with all else equal.[30] Absent the disruptions, inflation would have been 1.1 p.p. lower in December 2021 (6.1% rather than 7.2%), which is in line with estimates in the literature (e.g., di Giovanni et al., 2023, Comin et al., 2023).

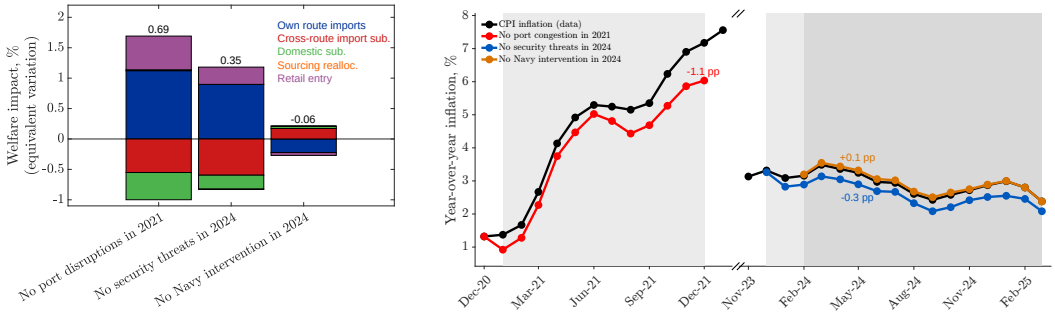
6.3. *Security threats in the Red Sea in 2023–2024*

Houthi forces began launching attacks on commercial vessels transiting the Red Sea and the Gulf of Aden in late 2023, forcing shipping companies to reroute vessels away from the Suez Canal to the Cape of Good Hope. We showed in Section 3 that this added roughly 600–800 nautical miles to rerouted voyages from Asia and 1,000–1,250 to those from the Middle East, or about 4 to 5 additional days of travel in our sample. Utilization on the routes from Asia to the East Coast via the Suez Canal and from the Middle East to the East Coast, z_{AES} and z_{ME}, fell by 21 and 31 p.p. between March 2023 and March 2024. The disruptions we measure occurred *with* the U.S. Navy's protection of commercial shipping from December 2023 onward (Section 7 quantifies the additional losses the intervention averted).

[29] This is computed by integrating the margin contributions along the path between the baseline and counterfactual economies, so the channel contributions sum to the total equivalent variation by construction.

[30] The counterfactual rate applies the model's deviation of the aggregate price level from the baseline, $P_t^{\text{total,Exp}}/P_t^{\text{total,Base}} - 1$, to observed CPI inflation. Although the model is not designed to reproduce CPI inflation, the exercise isolates the contribution of the disruption to measured inflation, holding everything else fixed.

FIGURE 6.3.—Welfare and price impact of shipping disruptions



(a) Welfare impact decomposition

(b) Price impact of disruptions

Note: Panel (a) shows the welfare impact of shipping disruptions, measured as in Equation (6.1), decomposed into the margins of Proposition 4.1. Panel (b) shows year-over-year CPI inflation (black) and the inflation rate implied by the model's deviation of the aggregate price index from the baseline. The horizontal axis is broken between January 2022 and November 2023. Shaded areas mark the windows over which welfare is evaluated. Labels report the gap between the counterfactual and actual inflation as of December 2021 and May 2024.

We now build a counterfactual economy in which utilization on both routes is held at its August 2023 values (0.857 and 0.809 for the Asia–East Coast via the Suez Canal and the Middle East–East Coast routes, more than 26 p.p. above their May 2024 troughs) from December 2023 to March 2025, the period during which utilization was significantly below trend and over which we evaluate welfare. As in Section 6.2, we solve each month for the exogenous components x_{AES} and x_{ME} that deliver these values, keeping all other parameters at their baseline values. Table 6.1 reports the effects at the May 2024 trough, and Figures A6 and A7 in Appendix C.5 show the key aggregates.

Without security threats, shipping firms reallocate capacity from other Asian routes back to the Suez route, while idle capacity remains about the same as in the baseline. The number of trips per ship on the Suez route from Asia rises from 1.6 to 2.7 per month. Congestion rises as capacity flows back into it, partly offsetting the utilization gain. A model with exogenous trade costs would overstate the benefit of removing the disruption by ignoring this response. Despite the shift away from the routes to the West Coast and via the Panama Canal, total U.S.-bound potential capacity rises only modestly, from 42 to 47 million metric tons in May 2024.

Imports rise with the additional supply. The import share of GDP in May 2024 increases by 0.3 p.p., driven by goods from Asia and the Middle East, whose retail prices fall by 9.6% and 24.6%, respectively. The retail prices of goods from Europe and South America are practically unchanged. The aggregate price index falls by 0.3% relative to the baseline.

Figure 6.3a reports the welfare impact. Consumers would be willing to pay 0.35% of output to avoid disruptions caused by security threats in the Red Sea. The decomposition shows that the losses were primarily driven by a fall in imports on the affected routes, with cross-route substitution absorbing most of the adjustment. The Red Sea attacks were less damaging than

the 2021 port disruptions, in part because U.S. naval forces helped contain them before they worsened. The value of that intervention is quantified in Section 7.

6.4. *Robustness*

We assess the sensitivity of the welfare estimates to the parameters governing the markups of shipping firms and importers (σ and θ), shipping costs (ζ), importers' sourcing decision (λ), consumer demand (ϱ and ρ), and to the choice of the pre-disruption anchor date. For each perturbation, we hold the model's fit to the data fixed, which requires recalibrating the internally determined parameters and simulating the baseline economy so that it continues to reproduce the targeted moments and time paths of Section 5.3.[31]

Table A4 in Appendix C.7 reports the equivalent variation for the 2021 West Coast port disruptions and the 2024 Red Sea security threats under $\pm 10\%$ perturbations to σ, θ, ζ, λ, ϱ and ρ, as well as for an alternative pre-disruption anchor date (December 2019 and July 2023). The welfare costs remain within 0.58–0.88% of output for the 2021 West Coast port disruptions and 0.27–0.59% for the 2024 Red Sea security threats, leaving the qualitative conclusions unchanged.

7. The value of Navy protection

Was the U.S. Navy's intervention in the Red Sea worth its cost? We infer the fiscal costs of the military operations from military records and use our model to estimate the economic losses averted. We find that the benefits to U.S. consumers exceeded the budgetary costs several times over—though, as we emphasize below, this comparison depends on assumptions about the counterfactual evolution of disruptions in the absence of the intervention.

In response to the Houthi attacks in the Red Sea, the U.S. Navy and other allies formed a coalition in December 2023 and launched *Operation Prosperity Guardian* to patrol the Red Sea and actively shoot down incoming missiles and drones. Although more than twenty nations eventually joined the coalition, the U.S. supplied the large majority of the naval assets and, together with the U.K., conducted most of the strike operations.[32]

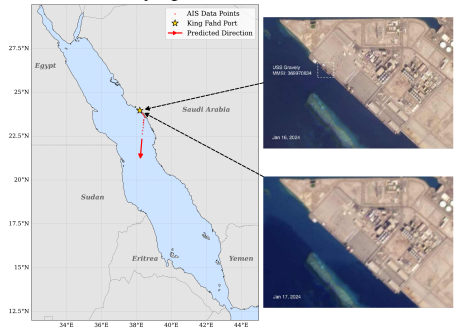
Following the attack on the *Maersk Hangzhou* on December 30, 2023, the U.S. Navy responded with the destroyers USS Laboon and USS Gravely as well as the aircraft carrier USS Dwight D. Eisenhower. The first airstrikes were launched against Houthi targets in Yemen on January 12, 2024. The campaign was followed by a decline in attacks on commercial vessels,

[31] The welfare estimates are invariant to the congestion externality parameters ε_k and ϕ by construction: the counterfactuals hold the utilization paths z_r at their pre-disruption levels so the congestion externality does not directly impact the results. It shapes the policy experiments in Section 8, where z_r is endogenous.

[32] The remaining partners contributed staff or vessels operating under separate mandates.

culminating in a ceasefire in May 2025. Figure 7.1 maps the USS Gravely's patrol movements from AIS data, illustrating the mobilized naval presence required to safeguard the corridor.[33]

FIGURE 7.1.—USS Gravely operations in the Red Sea (January 17, 2024)



Note: The figure illustrates the movement of the Arleigh Burke-class guided-missile destroyer USS Gravely (DDG 107) using AIS data. Because AIS transmissions for U.S. Navy ships are typically disabled during active military operations for security reasons, the available data are limited and primarily concentrated around periods when the ship was departing for operations or docking afterward.

We begin by estimating the fiscal cost of the U.S. Navy's intervention. Table 7.1 reports our estimate, built from naval operations and munitions expenditures documented in Navy and CBO budget justifications and other reporting on military usage. The naval operations totaled $1.3 to 3.0 billion, or 0.004 to 0.01% of GDP, for the seven naval assets deployed for about eight months between late 2023 and the summer of 2024 (the range depends on whether overhead and other indirect costs are included). The munitions used to counter the Houthi attacks amounted to an additional $1.2 to 2.4 billion, or 0.004 to 0.008% of GDP. The Israel Security Supplemental Appropriations Act of April 2024 included $2.44 billion to help address the additional expenses of the increased naval operations against the Houthi threats, which likely covered part of these outlays. Without supplemental funds, the total cost was between $2.4 and $5.5 billion, or 0.008 to 0.02% of U.S. GDP. With them, the total cost was between $4.9 and $7.9 billion, or 0.02 to 0.03% of GDP. These figures mostly cover operations through July 2024, so the fiscal cost over the full welfare window below is somewhat larger as operations continued into 2025.

We then use the model to estimate the value of the Navy intervention, complementing direct estimates of the welfare cost of maritime insecurity such as Besley et al. (2015) for Somali piracy. We construct three counterfactual economies in which shipping disruptions along the Asia–East Coast route via the Suez Canal and the Middle East–East Coast route continue to grow throughout 2024 and early 2025. Attacks on commercial vessels were escalating in the months before the coalition deployed, so utilization on these routes would plausibly have continued to deteriorate absent the intervention. The medium scenario extends the pre-intervention

[33] Additional satellite figures tracking USS Gravely are provided in Figure A3, Appendix A.

TABLE 7.1

ESTIMATED COST OF U.S. NAVY RED SEA OPERATIONS

	Quantity	Unit Cost (US$M)	Total Cost (US$M)	% of GDP	Source
		Panel A: Naval Operations			
Arleigh Burke-class destroyer	4 × 250 days	87–211	237–578	0.0008–0.0020	[1],[2],[3]
Carrier air wing	1 × 250 days	575–1,395	394–956	0.0013–0.0033	[1],[2],[3]
Nimitz-class aircraft carrier	1 × 250 days	752–1,830	515–1,253	0.0018–0.0043	[1],[2],[3]
Ticonderoga-class cruiser	1 × 250 days	187–380	128–260	0.0004–0.0009	[1],[2],[3]
Total			**1,273–3,048**	**0.0043–0.0104**	
		Panel B: Munitions			
SM-2	120	3.6–4.3	438–519	0.0015–0.0018	[4],[5]
SM-3/ESSM	20	1.9–25.3	39–507	0.0001–0.0017	[4],[5],[6]
SM-6	80	4.0–9.6	321–766	0.0011–0.0026	[5],[6]
Tomahawk cruise missile	135	2.1–3.9	287–531	0.0010–0.0018	[1],[7]
AMRAAM/AIM-9X	60	0.9–1.1	55–67	0.0002	[6],[8]
JDAM/AGM-114 Hellfire	420	0.05–0.12	20–49	0.0001–0.0002	[6],[8]
Hypervelocity Projectile	160	0.1	17	0.0001	[5],[9]
Total			**1,176–2,439**	**0.0040–0.0083**	
		Panel C: Other			
Israel Sec. Supp. (ISS) Approp. Act, 2024			2,440	0.0083	[10]
		Panel D: Total Cost			
Total (without ISS)			**2,449–5,487**	**0.0084–0.0187**	
Total (with ISS)			**4,889–7,927**	**0.0167–0.0271**	

Notes: The table provides estimates of the cost of U.S. Navy operations in the Red Sea between late 2023 and July 2024 (shares are calculated using 2024 nominal GDP of $29.3 trillion). *Panel A:* Naval operations include one Nimitz-class aircraft carrier (USS Dwight D. Eisenhower), four Arleigh Burke-class destroyers (USS Gravely, USS Mason, USS Laboon, and USS Carney), one carrier air wing (Carrier Air Wing 3), and one Ticonderoga-class cruiser (USS Philippine Sea) deployed for approximately 250 days between October 2023 and July 2024. The lower bound for unit cost refers to direct costs (which include the costs of the combat unit and the personnel assigned to the unit itself), and the upper bound adds overhead and other indirect costs (such as medical care, training, and administration).

Panel B: SM and ESSM stand for Standard Missiles and Evolved SeaSparrow Missiles; SM, ESSM, and Tomahawk cruise missiles are ship-launched surface-to-air missiles. AMRAAM and AIM-9X stand for Advanced Medium-Range Air-to-Air Missiles and Air Intercept Missile-9X; these are air-to-air missiles. JDAM stands for Joint Direct Attack Munition; JDAM and AGM-114 Hellfire are air-to-surface missiles. Hypervelocity Projectiles are versatile surface-to-air/surface missiles and correspond to rounds from a five-inch gun.

Panel C: The Israel Security Supplemental provided funding for the Central Command operations and helped address the additional expenses related to the increased U.S. naval operations against Houthi threats. The amount shown likely covered part of the expenses in Panels A and B.

Sources: [1] Department of the Navy (2024b); [2] Congressional Budget Office (2021); [3] Congressional Budget Office (2025); [4] Department of the Navy (2024a); [5] Ziezulewicz (2025); [6] Department of Defense (2024); [7] Department of the Navy (2023); [8] Trevithick (2024); [9] Congressional Research Service (2021); [10] House Committee on Appropriations (2024).

downward trend in utilization, with z_{AES} reaching 57% in November 2024; the small-shock scenario has the decline stabilize at that level earlier, in August 2024; and in the large-shock scenario the decline accelerates and persists until March 2025.[34] Figures A8 and A9 in Appendix C.6 display the utilization paths and other aggregates under the three scenarios.

Figure 6.3a shows the welfare impact had the attacks worsened and lasted longer. As consumer welfare is lower without the Navy's protection, consumers are worse off relative to the baseline economy. The value of the Navy's support between February 2024 and March 2025 amounted to 0.06% of output in the medium scenario, and to 0.04% and 0.08% in the small- and large-shock scenarios. In terms of prices, Figure 6.3b shows that inflation would have been marginally affected without the Navy intervention. These values are measured as the consumer-welfare equivalent variation and so exclude non-consumer surplus, such as the risk borne by crews and vessels. The magnitudes are nonetheless consistent with the modest but non-negligible share of U.S. imports transiting these routes.

[34] For the Middle East–East Coast route, utilization stabilizes at 45%, close to its February 2024 value. As before, the counterfactuals are computed by solving for x_{AES} and x_{ME} that deliver the assumed utilization paths.

To sum up, relative to the 0.35% of output welfare cost that materialized with the Navy in place (from Section 6.3), the intervention averted an additional 0.06% of output in losses in the medium scenario against a fiscal outlay of 0.008 to 0.03% of GDP, a benefit-cost ratio of roughly two to seven depending on the cost measure. The comparison is indicative rather than exact: the costs are U.S. outlays for a coalition effort and exclude partners' contributions and operations from late 2024 and early 2025, while the benefits are measured for U.S. consumers only, even though the corridor serves other countries' trade. Table A4 in Appendix C.7 reports the sensitivity of the medium-scenario welfare gain to alternative parameter values; across these perturbations, the gain ranges from 0.03% to 0.07% of output.

8. Tax policy in a congested shipping network

Motivated by the escalation of the trade war in 2025, we use the model to ask how tariffs and port fees operate when trade costs are the equilibrium outcome of a congested shipping network.[35] We now add a government that levies a region-specific tariff τ_o on importers and a port fee $\mathfrak{f}_o$ on shipping firms, and rebates the revenue lump sum to consumers.

The port fee raises the marginal cost of shipping along route r (Equation (4.8)):

$$mc_r = \frac{g(K_r)}{x_r}\left[\frac{\overline{p}_n}{1/\zeta}\right]^{1/\zeta}\left[\frac{(\overline{p}_o+\mathfrak{f}_o)}{1/\nu}\right]^{1/\nu}. \tag{8.1}$$

The tariff instead enters the importers' problem, affecting their sourcing decision o_r and their demand for shipping firm j's services, with Equation (4.17) now given by

$$\mathrm{y}_r(j) = o_r\left[\frac{\theta\chi_r^{\sigma}\, p_o y_r^{(1-\sigma/\theta)}}{(1+\tau_o)p_r(j)}\right]^{1/(1-\sigma)}\mathcal{I}. \tag{8.2}$$

The government rebates total revenue T lump sum to consumers:[36]

$$T = \sum_{o=\{A,E,M,S\}}\sum_{r=\mathcal{R}_o}\left[\int_0^N \mathfrak{f}_o(\ell_r(j)\kappa_r(j))^{\nu}dj + \tau_o\, o_r\mathcal{I}\int_0^N p_r(j)y_r(j)dj\right]. \tag{8.3}$$

[35] The average effective U.S. tariff rate reached nearly 18% in October 2025—the highest level since 1934. That month, the U.S. imposed fees on Chinese-built and Chinese-owned vessels calling at American ports ($18 per net ton, or $120 per container, per voyage for non-Chinese operators of Chinese-built ships and $50 per net ton per voyage for Chinese-operated ships). The fees were later suspended in November 2025.

[36] Adding these tax policy instruments also alters the labor market clearing condition, with the labor demanded by domestic firms now $l_d = 1 - \frac{1}{w}\int_0^N \sum_{r\in\mathcal{R}}\left\{\overline{p}_n n_r(j)^{\zeta} + \frac{\overline{p}_o}{\nu(\overline{p}_o+\mathfrak{f}_o)}\left[\zeta\overline{p}_n n_r(j)^{\zeta} - \mu_r(j)\right] + \overline{p}_{\kappa}\frac{\kappa_r(j)^2}{2}\right\}dj$.

Propositions B.3–B.5 in Appendix B.4 characterize the transmission of tariffs.[37] A tariff on region o operates through the conventional channel, raising importers' costs and the retail price of goods from that region, but it also reduces the capacity deployed on routes from region o, lowering congestion and the marginal cost of the shipping that remains. When the congestion externality is uninternalized, this second channel yields a first-order welfare benefit as in Section 6, with the tariff (rather than the disruption) withdrawing activity from saturated routes. Which channel dominates cannot be determined theoretically; it is a quantitative question.[38]

Table 8.1 compares the March 2025 baseline economy with an economy imposing an 18% tariff on Asian goods. To isolate the conventional channel, the first column holds route utilization, z_r, at its baseline path, so that congestion cannot respond to the reallocation of capacity; the second column allows the full response. The conventional channel alone reduces welfare by 0.287% of output: the retail price of Asian goods rises by 17.1%, the Asian import share falls by 19.5%, and the measure of importers falls by 11.3%. Allowing congestion to respond reverses the assessment. Capacity leaves the Asian routes and the price of shipping services falls by about 15.5% across all Asian routes as the tariff is now absorbed by cheaper shipping services rather than passed on to consumers. The retail price of Asian goods rises modestly (0.9%) and the Asian import share falls by only 2.9%. The tariff now yields a positive welfare effect (0.294% of output) as consumers are better off with cheaper imported goods.

TABLE 8.1

THE EFFECT OF AN 18% TARIFF ON ASIAN GOODS

	Baseline calibration (retail markup 25%)		Retail markup 50% ($\theta = 1/1.5$)	
	Conventional channel only	With decongestion response	Conventional channel only	With decongestion response
Shipping price, Asian routes (%)[a]	0.0, 0.0, 0.0	−15.6, −15.4, −15.9	0.0, 0.0, 0.0	−15.4, −15.1, −15.8
Retail price, Asian goods (%)	17.1	0.9	13.2	0.6
Aggregate price index (%)	0.4	−0.0	0.4	−0.0
Import share, Asia (%)	−19.5	−2.9	−14.9	−2.2
Import share, total (%)	−11.4	0.1	−9.5	0.1
Potential shipping capacity, total (%)	−20.8	−10.2	−19.1	−10.3
Effective shipping capacity, total (%)	−21.2	2.5	−19.4	2.6
Measure of importers (%)	−11.3	0.4	−9.4	0.4
Tariff revenue (% of GDP)	0.25	0.28	0.21	0.24
Welfare: equivalent variation (% of output)	−0.287	+0.294	−0.271	+0.250

Notes: All columns compare a March 2025 economy with an 18% tariff on goods from Asia, revenue rebated lump sum, against the corresponding no-tariff baseline; rows are percent deviations unless noted. In the "conventional channel only" columns, route-level shipping utilization is held at its baseline path, so congestion does not respond to the reallocation of capacity; the within-panel difference is the decongestion channel. The right panel recalibrates and simulates the economy to a 50% retail markup, with the internally determined parameters re-solved so the baseline reproduces the same targets. [a] The shipping prices are for the Asia–East Coast via the Panama Canal, the Suez Canal, and to the West Coast, respectively.

[37] Port fees levied on shipping firms operate through the same margins (see Propositions B.6–B.8 in Appendix B.5).

[38] The first-best instrument for the congestion externality would be Pigouvian, i.e., a tax on the capacity that shipping firms deploy on congested routes. The tariff instead acts as a second-best substitute, suppressing the demand that draws capacity onto congested routes (Bhagwati and Ramaswami, 1963, Brueckner, 2002).

We do not read the net gain as a robust policy prescription. Doubling the retail markup dampens the transmission of the tariff rather than reversing it. This is consistent with Proposition B.4, in which the response of retail output to tariffs scales with $\theta/(1-\theta)$, as both the conventional welfare loss and the decongestion welfare gain shrink in the right panel of Table 8.1. The columns of Table 8.1 bracket the welfare effect of the tariff. As the model abstracts from retaliation and terms-of-trade effects, the analysis isolates the congestion margin rather than evaluating the 2025 tariffs overall.

9. Improving port infrastructure

Recent port infrastructure investments have targeted reductions in turnaround times, the expansion of physical capacity, and better inland connectivity. These upgrades alleviate congestion and raise the utilization of existing infrastructure. In Appendix C.8 we evaluate such an improvement by lowering the congestion shifter ε_k by 10%, equivalent to a 10% reduction in anchorage waiting times. Consumers would pay 0.21% of output for the improvement.

10. Conclusion

We have used satellite vessel-tracking data to compute novel measures of potential and effective shipping capacity for the universe of U.S.-bound container ships from 2016 to 2025. This exercise shows that potential and effective capacity can diverge sharply amid shipping disruptions. During the 2021–2022 supply chain crisis, shipping capacity rose to record highs, yet nearly a third of it was absorbed by disruptions. The Houthi attacks in the Red Sea caused shipping utilization losses of similar magnitudes on the affected routes.

To quantify the macroeconomic impact of these shipping disruptions, we built a multi-route general equilibrium model of the shipping market, in which shipping firms reallocate a common fleet and importers adjust sourcing. Congestion externalities and rents from importers' sourcing imply that the reallocation of shipping capacity across routes has first-order welfare effects. Counterfactuals for the 2021 West Coast port crisis and the 2024 Red Sea attacks indicate sizeable welfare losses, about 0.69% and 0.35% of output, respectively.

The framework enables the evaluation of policies whose effects run through the shipping market. Navy protection in the Red Sea generates benefits that surpass its fiscal cost; an 18% tariff on Asian goods decongests the routes it taxes, enough to reverse its conventional welfare cost; and port infrastructure investments that reduce port delays yield gains through the same congestion margin. Promising extensions include endogenizing shipbuilding and allowing for strategic interactions among shipping alliances. We leave these extensions for future work.

REFERENCES

D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016, 2012.

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

T. Allen and C. Arkolakis. The Welfare Effects of Transportation Infrastructure Improvements. *The Review of Economic Studies*, 89(6):2911–2957, 2022.

American Society of Civil Engineers. 2025 Infrastructure Report Card: Ports, 2025.

M. Amiti, S. J. Redding, and D. E. Weinstein. The Impact of the 2018 Tariffs on Prices and Welfare. *Journal of Economic Perspectives*, 33(4):187–210, 2019.

J. E. Anderson and D. Marcouiller. Insecurity and the Pattern of Trade: An Empirical Investigation. *The Review of Economics and Statistics*, 84(2):342–352, 2002.

C. Arkolakis, A. Costinot, and A. Rodríguez-Clare. New Trade Models, Same Old Gains? *American Economic Review*, 102(1):94–130, 2012.

J. Asturias. Endogenous Transportation Costs. *European Economic Review*, 123:103366, 2020.

X. Bai, Z. Ma, Y. Hou, Y. Li, and D. Yang. A Data-Driven Iterative Multi-Attribute Clustering Algorithm and Its Application in Port Congestion Estimation. *IEEE Transactions on Intelligent Transportation Systems*, 24(11): 12026–12037, 2023.

X. Bai, J. Fernández-Villaverde, Y. Li, and F. Zanetti. The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory. *American Economic Review*, forthcoming, 2026a.

X. Bai, J. Fernández-Villaverde, Y. Li, and F. Zanetti. State Dependence of Monetary Policy During Global Supply Chain Disruptions. *Journal of Political Economy: Macroeconomics*, forthcoming, 2026b.

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

D. R. Baqaee and E. Farhi. Productivity and Misallocation in General Equilibrium. *The Quarterly Journal of Economics*, 135(1):105–163, 2020.

T. Besley, T. Fetzer, and H. Gould. The Welfare Cost of Lawlessness: Evidence From Somali Piracy. *Journal of the European Economic Association*, 13(2):203–239, 2015.

J. Bhagwati and V. K. Ramaswami. Domestic Distortions, Tariffs and the Theory of Optimum Subsidy. *Journal of Political Economy*, 71(1):44–50, 1963.

J. Blaum, F. Esposito, and S. Heise. Input Sourcing Under Climate Risk: Evidence From U.S. Manufacturing Firms. *Review of Economic Studies*, forthcoming, 2026.

C. E. Boehm, A. Flaaen, and N. Pandalai-Nayar. Input Linkages and the Transmission of Shocks: Firm-Level Evidence From the 2011 Tohoku Earthquake. *The Review of Economics and Statistics*, 101(1):60–75, 2019.

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

G. Brancaccio, M. Kalouptsidi, and T. Papageorgiou. Rigidities in Transportation and Supply Chain Disruptions. *AEA Papers and Proceedings*, 115:543–550, 2025.

G. Brancaccio, M. Kalouptsidi, and T. Papageorgiou. Investment in Infrastructure and Trade: The Case of Ports. *Journal of Political Economy*, forthcoming, 2026.

J. K. Brueckner. Airport Congestion When Carriers Have Market Power. *American Economic Review*, 92(5): 1357–1375, 2002.

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

L. Caliendo and F. Parro. Estimates of the Trade and Welfare Effects of NAFTA. *The Review of Economic Studies*, 82(1):1–44, 2015.

L. Caliendo, M. Dvorkin, and F. Parro. Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock. *Econometrica*, 87(3):741–835, 2019.

Y. Carrière-Swallow, P. Deb, D. Furceri, D. Jiménez, and J. D. Ostry. Shipping Costs and Inflation. *Journal of International Money and Finance*, 130:102771, 2023.

V. M. Carvalho, M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi. Supply Chain Disruptions: Evidence From the Great East Japan Earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321, 2021.

A. Cavallo, G. Gopinath, B. Neiman, and J. Tang. Tariff Pass-Through at the Border and at the Store: Evidence From U.S. Trade Policy. *American Economic Review: Insights*, 3(1):19–34, 2021.

Clarksons Research Services Limited. Shipping Intelligence Network Timeseries, 2025. SCFI series 98514 and 98515; underlying source: Shanghai Shipping Exchange.

A. K. Coşar, S. Osotimehin, and L. Popov. The Long-Run Effects of Transportation Productivity on the US Economy. Working Paper 33248, NBER, 2024.

D. A. Comin, R. C. Johnson, and C. J. Jones. Supply Chain Constraints and Inflation. Working Paper 31179, NBER, 2023.

Congressional Budget Office. The U.S. Military's Force Structure: A Primer, 2021 Update, 2021.

Congressional Budget Office. The U.S. Military's Force Structure: Fiscal Year 2025 Update to Personnel Numbers and Costs, 2025.

Congressional Research Service. Navy Lasers, Railgun, and Gun-Launched Guided Projectile: Background and Issues for Congress, 2021.

B. De Borger, S. Proost, and K. Van Dender. Port Activities, Hinterland Congestion, and Optimal Government Policies: The Role of Vertical Integration in Logistic Operations. *Journal of Transport Economics and Policy*, 45 (2):247–275, 2011.

Department of Defense. FY 2025 Program Acquisition Costs by Weapon System. Budget materials, Office of the Under Secretary of Defense (Comptroller), 2024.

Department of the Navy. FY 2024 Weapons Procurement, 2023.

Department of the Navy. FY 2025 Weapons Procurement, 2024a.

Department of the Navy. Unprecedented: Dwight D. Eisenhower Carrier Strike Group Returns From Combat Deployment. Press Release, 2024b.

J. di Giovanni, Ş. Kalemli-Özcan, A. Silva, and M. A. Yıldırım. Quantifying the Inflationary Impact of Fiscal Stimulus Under Supply Constraints. Working Paper 30892, NBER, 2023.

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

P. D. Fajgelbaum and E. Schaal. Optimal Transport Networks in Spatial Equilibrium. *Econometrica*, 88(4): 1411–1452, 2020.

P. D. Fajgelbaum, P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal. The Return to Protectionism. *The Quarterly Journal of Economics*, 135(1):1–55, 2020.

R. C. Feenstra, P. Luck, M. Obstfeld, and K. N. Russ. In Search of the Armington Elasticity. *The Review of Economics and Statistics*, 100(1):135–150, 2018.

J. Fernández-Villaverde, Y. Li, L. Xu, and F. Zanetti. Charting the Uncharted: The (Un)Intended Consequences of Oil Sanctions and Dark Shipping. Working Paper 33486, NBER, 2025.

S. Ganapati, W. F. Wong, and O. Ziv. Entrepôt: Hubs, Scale, and Trade Costs. *American Economic Journal: Macroeconomics*, 16(4):239–278, 2024.

R. Glick and A. M. Taylor. Collateral Damage: Trade Disruption and the Economic Impact of War. *The Review of Economics and Statistics*, 92(1):102–127, 2010.

R. Greenwood and S. G. Hanson. Waves in Ship Prices and Investment. *The Quarterly Journal of Economics*, 130 (1):55–109, 2015.

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

House Committee on Appropriations. House Passes Series of Security Supplemental Bills. Press Release, 2024.

C. R. Hulten. Growth Accounting With Intermediate Inputs. *The Review of Economic Studies*, 45(3):511–518, 1978.

D. Hummels, V. Lugovskyy, and A. Skiba. The Trade Reducing Effects of Market Power in International Shipping. *Journal of Development Economics*, 89(1):84–97, 2009.

M. Isaacson and H. Rubinton. Shipping Prices and Import Price Inflation. *Federal Reserve Bank of St. Louis Review*, 105(2):89–107, 2023.

M. Kalouptsidi. Time to Build and Fluctuations in Bulk Shipping. *American Economic Review*, 104(2):564–608, 2014.

J. LaBelle and A. M. Santacreu. Global Supply Chain Disruptions and Inflation During the COVID-19 Pandemic. *Federal Reserve Bank of St. Louis Review*, 104(2):78–91, 2022.

N. Limão and A. J. Venables. Infrastructure, Geographical Disadvantage, Transport Costs, and Trade. *The World Bank Economic Review*, 15(3):451–479, 2001.

V. Lugovskyy, A. Skiba, and D. Terner. Unintended Consequences of Environmental Regulation of Maritime Shipping: Carbon Leakage to Air Shipping. *Journal of International Economics*, 155:104081, 2025.

R. Marto. The Service Economy and the Rise in Markups. *Mimeo*, 2025.

M. J. Melitz and S. J. Redding. New Trade Models, New Welfare Implications. *American Economic Review*, 105(3): 1105–1146, 2015.

T. Notteboom, A. Pallis, and J.-P. Rodrigue. *Port Economics, Management and Policy*. Routledge, London, 2021.

T. Notteboom, H. Haralambides, and K. Cullinane. The Red Sea Crisis: Ramifications for Vessel Operations, Shipping Networks, and Maritime Supply Chains. *Maritime Economics & Logistics*, 26(1):1–20, 2024.

M. Stopford. *Maritime Economics*. Routledge, London, third edition, 2008.

S. P. Strandenes and P. B. Marlow. Port Pricing and Competitiveness in Short Sea Shipping. *International Journal of Transport Economics*, 27(3):315–334, 2000.

The White House. Fact Sheet: The Bipartisan Infrastructure Deal. Statements and Releases, November 6, 2021.

J. Trevithick. 770 Weapons Expended by Eisenhower Carrier Strike Group on Historic Red Sea Deployment. The War Zone, 2024.

UNCTAD. Review of Maritime Transport 2022, 2022.

UNCTAD. Review of Maritime Transport 2024, 2024.

W. F. Wong. The Round Trip Effect: Endogenous Transport Costs and International Trade. *American Economic Journal: Applied Economics*, 14(4):127–166, 2022.

W. Y. Yap and D. Yang. Geopolitical Tension and Shipping Network Disruption: Analysis of the Red Sea Crisis on Container Port Calls. *Journal of Transport Geography*, 121:104004, 2024.

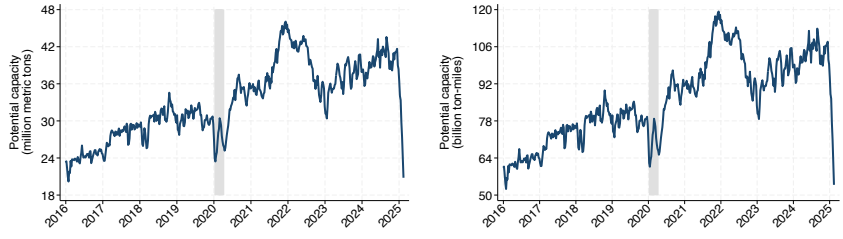
G. Ziezulewicz. Navy Just Revealed Tally of Surface-to-Air Missiles Fired in Ongoing Red Sea Fight. The War Zone, 2025.

ONLINE APPENDIX

A: ADDITIONAL EMPIRICAL RESULTS

This appendix collects additional empirical results referenced in the main text.

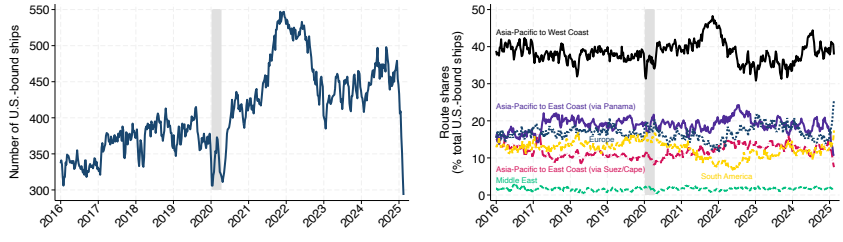
FIGURE A1.—Potential capacity from unique container ships on U.S.-bound routes, 2016–2025



(a) Potential capacity (tons) (b) Potential capacity (ton-miles)

Note: Panel A1a shows the aggregate potential shipping capacity in million metric tons and Panel A1b in billion ton-miles. Estimates are derived from the set of unique container ships operating on U.S.-bound routes. Shaded areas indicate recessions.

FIGURE A2.—U.S.-bound container ships, 2016–2025



(a) Total number (b) Route shares

Note: Panel A2a plots the total number of container ships operating on U.S.-bound routes. Panel A2b shows the share of ships by route. Shaded areas indicate recessions.

FIGURE A3.—USS Gravely near NSA Bahrain and the Port of Duqm, Oman



(a) Bahrain, December 13, 2023 (b) Oman, March 8, 2024

Note: Movements of the Arleigh Burke-class guided-missile destroyer USS Gravely (DDG 107) near Naval Support Activity (NSA) Bahrain (Panel A3a) and near the Port of Duqm in Oman (Panel A3b) using AIS data.

B: THEORY

This appendix proves the propositions of Section 4, characterizes the constrained social planner's problem, and states the tariff and port-fee counterparts of Proposition 4.1 used in Section 8.

B.1. *Proof of Proposition 4.1 (Impact of shipping disruptions on real aggregate output)*

In this economy, real aggregate output is equivalent to consumer welfare and thus $Y \equiv \left[\gamma_d d^{\varrho} + (1-\gamma_d)c^{\varrho}\right]^{1/\varrho}$. Applying the chain rule, the elasticity of aggregate output with respect to route-r aggregate utilization, $(\partial Y/\partial X_r)\,(X_r/Y)$, where the route-$r$ aggregate utilization is defined as $X_r \equiv \left[\int_0^N x_r(j)^{\sigma/(1-\sigma)}dj\right]^{(1-\sigma)/\sigma}$, is given by

$$\frac{\partial Y}{\partial X_r}\frac{X_r}{Y} = \left(\frac{\partial Y}{\partial d}\frac{d}{Y}\right)\left(\frac{\partial d}{\partial X_r}\frac{X_r}{d}\right) + \left(\frac{\partial Y}{\partial c}\frac{c}{Y}\right)\left\{\sum_{o\in\{A,E,M,S\}}\left(\frac{\partial c}{\partial c_o}\frac{c_o}{c}\right)\right. \tag{B.1}$$

$$\left.\times\sum_{h|o}\omega_h\left[\left(\frac{\partial y_h}{\partial X_r}\frac{X_r}{y_h}\right) + \left(\frac{\partial o_h}{\partial X_r}\frac{X_r}{o_h}\right) + \left(\frac{\partial \mathcal{I}}{\partial X_r}\frac{X_r}{\mathcal{I}}\right)\right]\right\},$$

where $\omega_h \equiv o_h y_h \mathcal{I}/c_o$ is the share of route h goods in the consumption of goods from region o. As there is only one route for goods from Europe, the Middle East, and South America to arrive in the U.S., $\omega_h = 1$ for $o \in \{E, M, S\}$.

From the consumer's first-order conditions, it is straightforward to deduce that

$$\frac{\partial Y}{\partial d}\frac{d}{Y} = \frac{p_d d}{w+\Lambda} \equiv \text{Domestic share in aggregate output},$$

$$\frac{\partial Y}{\partial c}\frac{c}{Y} = \frac{P^{\text{imp}} c}{w+\Lambda} \equiv \text{Total import share in aggregate output},$$

and

$$\frac{\partial c}{\partial c_o}\frac{c_o}{c} = \frac{p_o c_o}{P^{\text{imp}} c} \equiv \text{Region } o\text{'s import share of total imports},$$

where $P^{\text{imp}} \equiv \left(\sum_{o\in\mathcal{O}} \gamma_o^{1/(1-\rho)} p_o^{\rho/(\rho-1)}\right)^{(\rho-1)/\rho}$ is the imported goods price index. Route h's Domar weight is then simply given by

$$\text{Route } h\text{'s share in aggregate output} \equiv \frac{\partial Y}{\partial c}\frac{c}{Y} \times \frac{\partial c}{\partial c_o}\frac{c_o}{c} \times \frac{o_h y_h \mathcal{I}}{c_o}.$$

The impact of shipping disruptions along route r on the consumption of and spending on domestic goods depends on how the shipping disruptions affect the labor demanded by producers of domestic goods through general equilibrium effects. This in turn depends on how shipping profits and entry are affected by the disruptions. Similarly, the impact of shipping disruptions on the measure of importers depends on the free-entry condition in the retail sector.

B.2. *Impact of shipping disruptions on importers' output and sourcing*

B.2.1. *Proposition B.1 (Impact of shipping disruptions on importers' output and sales)*

PROPOSITION B.1: (IMPACT OF SHIPPING DISRUPTIONS ON IMPORTERS' OUTPUT AND SALES) *Ceteris paribus*, a 1% increase in route-r utilization, X_r, raises importers' output and sales from that same route by

$$\frac{\partial p_o y_r}{\partial X_r}\frac{X_r}{p_o y_r} = \frac{\partial y_r}{\partial X_r}\frac{X_r}{y_r} \approx \frac{\theta}{1-\theta} > 0.$$

In contrast, *ceteris paribus*, a 1% increase in route-r utilization has no effect on importers' output and sales from other routes, i.e.,

$$\frac{\partial p_{\hat{o}} y_h}{\partial X_r}\frac{X_r}{p_{\hat{o}} y_h} = \frac{\partial y_h}{\partial X_r}\frac{X_r}{y_h} \approx 0.$$

PROOF: Start with importers' demand for the shipping service provided by firm j along route r conditional on their optimal choice of specialization

$$y_r(j) = \left[\frac{\sigma\theta\chi_r^{\sigma} p_o y_r^{(1-\sigma/\theta)} x_r(j)}{g(K_r)\left(\frac{\overline{p}_n}{1/\zeta}\right)^{1/\zeta}\left(\frac{\overline{p}_o}{1/\nu}\right)^{1/\nu}}\right]^{1/(1-\sigma)},$$

where the price of the shipping services was replaced by the firm's markup and marginal cost and, for simplicity, an interior solution for the shipping firm's optimal loading was assumed. Exponentiate both sides, aggregate over all shipping firms, and rearrange to obtain importers' supply conditional on their optimal choice of specialization, or

$$y_r = \left[\frac{\sigma\theta\chi_r p_o X_r}{g(K_r)\left(\frac{\overline{p}_n}{1/\zeta}\right)^{1/\zeta}\left(\frac{\overline{p}_o}{1/\nu}\right)^{1/\nu}}\right]^{\theta/(1-\theta)}$$

as well as their sales

$$p_o y_r = \left[\frac{\sigma\theta\chi_r p_o^{1/\theta} X_r}{g(K_r)\left(\frac{\overline{p}_n}{1/\zeta}\right)^{1/\zeta}\left(\frac{\overline{p}_o}{1/\nu}\right)^{1/\nu}} \right]^{\theta/(1-\theta)}.$$

The elasticity of retail output to shipping disruptions is

$$\frac{\partial y_h}{\partial X_r}\frac{X_r}{y_h} = \begin{cases} \frac{\theta}{1-\theta}\left[\underbrace{1}_{\text{PE effect}} + \underbrace{\frac{\partial p_o}{\partial X_r}\frac{X_r}{p_o} - \frac{\partial g(K_h)}{\partial X_r}\frac{X_r}{g(K_h)}}_{\text{GE effect}}\right] & \text{if route } h = r \\ \frac{\theta}{1-\theta}\left[\underbrace{0}_{\text{PE effect}} + \underbrace{\frac{\partial p_{\hat{o}}}{\partial X_r}\frac{X_r}{p_{\hat{o}}} - \frac{\partial g(K_h)}{\partial X_r}\frac{X_r}{g(K_h)}}_{\text{GE effect}}\right] & \text{if route } h \neq r, \end{cases}$$

where $p_{\hat{o}}$ is the retail price of the goods associated with route h when the region of origin is not o. The elasticity of retail sales to shipping disruptions is

$$\frac{\partial p_{\tilde{o}} y_h}{\partial X_r}\frac{X_r}{p_{\tilde{o}} y_h} = \begin{cases} \frac{\theta}{1-\theta}\left[\underbrace{1}_{\text{PE effect}} + \underbrace{\frac{1}{\theta}\frac{\partial p_o}{\partial X_r}\frac{X_r}{p_o} - \frac{\partial g(K_h)}{\partial X_r}\frac{X_r}{g(K_h)}}_{\text{GE effect}}\right] & \text{if route } h = r \text{ and } p_{\tilde{o}} = p_o \\ \frac{\theta}{1-\theta}\left[\underbrace{0}_{\text{PE effect}} + \underbrace{\frac{1}{\theta}\frac{\partial p_{\hat{o}}}{\partial X_r}\frac{X_r}{p_{\hat{o}}} - \frac{\partial g(K_h)}{\partial X_r}\frac{X_r}{g(K_h)}}_{\text{GE effect}}\right] & \text{if route } h \neq r \text{ and } p_{\tilde{o}} = p_{\hat{o}}. \end{cases}$$

Q.E.D.

B.2.2. *Proposition B.2 (Impact of shipping disruptions on retail sourcing)*

PROPOSITION B.2: (IMPACT OF SHIPPING DISRUPTIONS ON RETAIL SOURCING) *Ceteris paribus*, a 1% increase in route-r utilization raises the share of importers sourcing goods from that same route by

$$\frac{\partial o_r}{\partial X_r}\frac{X_r}{o_r} \approx (1-o_r)\lambda\frac{\theta}{1-\theta}(\Pi_r + f) > 0,$$

where importers' profit from selling goods transported along route r (net of the fixed cost) is $\Pi_r + f = (1-\theta)\left[\frac{\sigma\theta\chi_r p_o^{1/\theta} X_r}{g(K_r)\left(\frac{\overline{p}_n}{1/\zeta}\right)^{1/\zeta}\left(\frac{\overline{p}_o}{1/\nu}\right)^{1/\nu}}\right]^{\theta/(1-\theta)}$.

Ceteris paribus, a 1% increase in route-r utilization decreases the share of importers sourcing goods from other routes according to

$$\frac{\partial o_h}{\partial X_r}\frac{X_r}{o_h} \approx -o_r\lambda\frac{\theta}{1-\theta}(\Pi_r + f) < 0,$$

where the optimal sourcing choice, o_r, is given by Equation (4.16).

PROOF: Using importers' first-order conditions, the profit from producing goods from route r conditional on their sourcing decision is

$$\Pi_r = (1-\theta)\left[\frac{\sigma\theta\chi_r p_o^{1/\theta} X_r}{g(K_r)\left(\frac{\overline{p}_n}{1/\zeta}\right)^{1/\zeta}\left(\frac{\overline{p}_o}{1/\nu}\right)^{1/\nu}}\right]^{\theta/(1-\theta)} - f.$$

The impact of route-r shipping disruptions on importers' sourcing choice of goods transported along route h is given by

$$\begin{aligned}\frac{\partial o_h}{\partial X_r}\frac{X_r}{o_h} &= \lambda X_r\left(\frac{\partial \Pi_h}{\partial X_r} - \sum_{h'} o_{h'}\frac{\partial \Pi_{h'}}{\partial X_r}\right)\\ &= \lambda(1-\theta)X_r\left(\frac{\partial p_{\tilde{o}} y_h}{\partial X_r} - \sum_{h'} o_{h'}\frac{\partial p_{\tilde{o}} y_{h'}}{\partial X_r}\right),\end{aligned}$$

where $p_{\tilde{o}}$ denotes the retail price of the goods sourced from the respective route. *Ceteris paribus,* the first-order effect of shipping disruptions on the share of importers sourcing goods from route r is

$$\frac{\partial o_r}{\partial X_r}\frac{X_r}{o_r} \approx (1-o_r)\lambda\theta p_o y_r = (1-o_r)\lambda\frac{\theta}{1-\theta}(\Pi_r + f),$$

and the first-order effect of shipping disruptions on the share of importers sourcing goods from routes $h \neq r$ is given by

$$\frac{\partial o_h}{\partial X_r}\frac{X_r}{o_h} = -o_r\lambda\theta p_o y_r = -o_r\lambda\frac{\theta}{1-\theta}(\Pi_r + f),$$

where o_r is given by Equation (4.16). Note that the weighted sum of these elasticities is such that $\sum_{h'} o_{h'}\left(\frac{\partial o_{h'}}{\partial X_r}\frac{X_r}{o_{h'}}\right) = 0$.

Q.E.D.

B.3. *The constrained social planner's problem*

The planner's problem can be solved in four stages: (1) an inner cost-minimization problem over speed and loading per route given capacity; (2) the social price of shipping per route that takes into account that importers have a downward-sloping demand curve and that they can reallocate their sourcing decisions across routes; (3) a capacity allocation per route that internalizes the congestion effect that additional capacity generates; and (4) the characterization of the social marginal values as Lagrange multipliers emerging from the optimization such that all markets clear.

B.3.1. *Stage 1: Per-route shipping cost minimization*

For each route r and given capacity K_r, the choice of average speed s_r and loading ℓ_r is a self-contained cost-minimization problem. It involves only the route's own shipping technology and own input prices, and thus does not interact with any other route's choices.

In Stage 1, the planner solves a cost-minimization problem over these margins for each route, subject to the aggregate shipping technology $Q_r \equiv x_r\, A_r\, s_r^{\alpha}\, \ell_r\, K_r \,/\, g(K_r)$ and taking K_r as given. The solution to the shipping cost-minimization problem yields the same marginal cost of shipping as in Equation (4.8) as a function of K_r.[39]

B.3.2. *Stage 2: The social price of shipping*

The planner sets the route-level social price of shipping v_r^{social} for each $r \in \mathcal{R}$, taking as given importers' demand for shipping services and their sourcing decisions. Given that there are redeployment costs, the planner sets the social price of shipping as a markup over the marginal cost as in Equation (4.10) to cover these resources. However, the price elasticity of demand differs from the decentralized equilibrium as the planner takes into account the price impact on importers' sourcing decisions, and the relevant elasticity is not the firm's inner elasticity.

Denote importers' total demand for shipping services from route r conditional on sourcing goods from that route by $Y_r^{\text{input}} \equiv N y_r(j)$, so that the aggregate shipping output is equal to the aggregate importer demand in line with $Q_r = Y_r^{\text{input}} o_r \mathcal{I}$. The output from that route (Equation (4.13)) can now be expressed as $y_r = \left[\chi_r N^{(1-\sigma)/\sigma} Y_r^{\text{input}}\right]^{\theta}$ and the optimal input demand that maximizes social profits (the analog to Equation (4.14)) is $Y_r^{\text{input}} = \left[\frac{\theta \chi_r^{\theta}\, v_o^{c,\text{social}} N^{(1-\sigma)\theta/\sigma}}{v_r^{\text{social}}}\right]^{1/(1-\theta)}$, where $v_o^{c,\text{social}}$ is the social marginal value of consumption from region o.

The price elasticity of demand now depends on how total demand for shipping services from route r, Y_r^{input}, responds to changes in the social price of shipping and how importers

[39] The corner case for loading $\ell_r = 1$ is handled by a multiplier $\mu_r > 0$ on the loading bound, with the marginal cost modified accordingly (cf. Footnote 19).

reallocate across routes when that price changes. In the decentralized equilibrium, the price elasticity of demand depends solely on the inner elasticity of substitution between shipping firms (and is equivalent to $1/(1-\sigma)$). In the planner's problem, the price elasticity of demand is $\xi_r = \frac{1}{1-\theta}\left\{1 - \lambda(1-o_r)\left[(1-\theta)v_r^{\text{social}}Y_r^{\text{input}} - \theta f\right]\right\}$. The relevant elasticity is now the elasticity of retail output to the total demand for imports, θ, not the firm-specific σ. It also takes into account that importers can reallocate their sourcing to other goods. The resulting social markup is

$$m_r^{\text{social}} = \frac{1}{\theta}\left\{\frac{1-\lambda(1-o_r)\left[(1-\theta)v_r^{\text{social}}Y_r^{\text{input}} - \theta f\right]}{1-\lambda(1-o_r)/\theta\left[(1-\theta)v_r^{\text{social}}Y_r^{\text{input}} - \theta f\right]}\right\}, \tag{B.2}$$

which could be greater or smaller than $1/\theta$ and the decentralized markup $1/\sigma$. The social markup is now a fixed point as it enters both the left- and right-hand sides.

B.3.3. *Stage 3: Capacity allocation*

As the planner sets the social price of shipping services as a markup over the marginal cost, the planner has resources to cover the aggregate capacity adjustment costs. When choosing where to allocate capacity, the planner takes congestion, $g(K_r)$, and the aggregate capacity constraint, $\sum_{r\in\mathcal{R}} K_r \leq K$, into account. The more capacity is allocated to a route, the higher is the marginal cost of shipping.

The first-order condition of the planner's problem for route-r total capacity, K_r, when the optimal loading of ships is interior (i.e., $\ell_r < 1$), yields a similar expression to Equation (4.12)

$$N\left\{\underbrace{\frac{\left(v_r^{\text{social}} - (1+\phi)mc_r\right)x_r n_r(s_r^\star)\ell_r^\star}{g(K_r)}}_{\text{profit per ton traveled along route } r} - \underbrace{\frac{1}{|\mathcal{R}|}\sum_h \frac{\left(v_h^{\text{social}} - (1+\phi)mc_h\right)x_h n_h(s_h^\star)\ell_h^\star}{g(K_h)}}_{\text{average profit per ton traveled}}\right\} = \overline{p}_\kappa\left[K_r - \underbrace{\frac{1}{|\mathcal{R}|}(K-\overline{K})}_{\text{average capacity per route}}\right], \tag{B.3}$$

which equates the marginal benefit of reallocating capacity to route r to its marginal cost. The right-hand side of Equation (B.3) is the same as each firm's (Equation (4.12)) summed over all N firms. The left-hand side differs in two important ways. First, the congestion effect is taken into account, which scales the marginal cost of shipping by $(1+\phi)$. Second, the social price of shipping includes a markup that takes into account how importers reallocate across routes and how total demand for imported goods changes when the social price of shipping changes.

B.3.4. *Stage 4: The social price of consumption and labor*

The social marginal values of consumption, $v_o^{c,\text{social}}$, are now obtained from the Lagrange multipliers on the imported goods market clearing condition for each region o, so that $v_o^{c,\text{social}} = \partial u/\partial c_o$. Similarly, the social price of domestically-produced goods is $v_d^{\text{social}} = \partial u/\partial d$. The multiplier on the labor resource constraint is fixed at 1 by the numeraire normalization (equivalent to normalizing $w = 1$), which implies that $v_d^{\text{social}} = 1/z_d$.

Finally, the social planner ensures that the labor resource constraint holds, so that $l_d + \sum_{r\in\mathcal{R}} \{mc_r Q_r - N\mu_r^{\text{loading}}/\nu + \overline{p}_\kappa K_r^2/(2N)\} = 1$. This guarantees that all resources at the disposal of the social planner, $1 + f\mathcal{I}$, are allocated efficiently.

B.4. *The impact of tariffs on real aggregate output*

PROPOSITION B.3: (IMPACT OF TARIFFS ON REAL AGGREGATE OUTPUT) The elasticity of real aggregate output Y to tariffs τ_o on goods from region $o \in \mathcal{O}$ is given by

$$\frac{\partial Y}{\partial \tau_o}\frac{\tau_o}{Y} = \sum_{h\in\mathcal{R}} \left(\begin{smallmatrix}\text{Route } h\text{'s share in}\\ \text{aggregate output}\end{smallmatrix}\right) \times \left(\frac{\partial y_h}{\partial \tau_o}\frac{\tau_o}{y_h} + \frac{\partial o_h}{\partial \tau_o}\frac{\tau_o}{o_h}\right)$$
$$+ \left(\begin{smallmatrix}\text{Import share in}\\ \text{aggregate output}\end{smallmatrix}\right) \times \left(\frac{\partial \mathcal{I}}{\partial \tau_o}\frac{\tau_o}{\mathcal{I}}\right) + \left(\begin{smallmatrix}\text{Domestic share in}\\ \text{aggregate output}\end{smallmatrix}\right) \times \left(\frac{\partial d}{\partial \tau_o}\frac{\tau_o}{d}\right),$$

where $\left(\begin{smallmatrix}\text{Route } h\text{'s share in}\\ \text{aggregate output}\end{smallmatrix}\right) = p_{o(h)} o_h y_h \mathcal{I}/(w + \Lambda + T)$ (with $o(h)$ route h's origin region), $\left(\begin{smallmatrix}\text{Domestic share in}\\ \text{aggregate output}\end{smallmatrix}\right) = p_d d/(w + \Lambda + T)$, and $\left(\begin{smallmatrix}\text{Import share in}\\ \text{aggregate output}\end{smallmatrix}\right) = 1 - \left(\begin{smallmatrix}\text{Domestic share in}\\ \text{aggregate output}\end{smallmatrix}\right)$.

PROOF: The result follows from the same steps as the proof of Proposition 4.1. *Q.E.D.*

Propositions B.4 and B.5 below decompose the route-specific elasticities $\frac{\partial y_h}{\partial \tau_o}\frac{\tau_o}{y_h}$ and $\frac{\partial o_h}{\partial \tau_o}\frac{\tau_o}{o_h}$.

PROPOSITION B.4: (IMPACT OF TARIFFS ON IMPORTERS' OUTPUT AND SALES) The elasticity of retail output to tariffs is

$$\frac{\partial y_h}{\partial \tau_o}\frac{\tau_o}{y_h} = \begin{cases} -\frac{\theta}{1-\theta}\left[\underbrace{\frac{\tau_o}{1+\tau_o}}_{\text{PE effect}} + \underbrace{\frac{\partial g(K_h)}{\partial \tau_o}\frac{\tau_o}{g(K_h)} - \frac{\partial p_o}{\partial \tau_o}\frac{\tau_o}{p_o}}_{\text{GE effect}}\right] & \text{if route } h\text{'s origin} = o \\ -\frac{\theta}{1-\theta}\left[\underbrace{0}_{\text{PE effect}} + \underbrace{\frac{\partial g(K_h)}{\partial \tau_o}\frac{\tau_o}{g(K_h)} - \frac{\partial p_{\hat{o}}}{\partial \tau_o}\frac{\tau_o}{p_{\hat{o}}}}_{\text{GE effect}}\right] & \text{if route } h\text{'s origin} \neq o, \end{cases}$$

where $p_{\hat{o}}$ is the retail price of the goods associated with route h when the region of origin is not o. The elasticity of retail sales to tariffs is

$$\frac{\partial p_{\tilde{o}} y_h}{\partial \tau_o}\frac{\tau_o}{p_{\tilde{o}} y_h} = \begin{cases} -\frac{\theta}{1-\theta}\left[\underbrace{\frac{\tau_o}{1+\tau_o}}_{\text{PE effect}} + \underbrace{\frac{\partial g(K_h)}{\partial \tau_o}\frac{\tau_o}{g(K_h)} - \frac{1}{\theta}\frac{\partial p_o}{\partial \tau_o}\frac{\tau_o}{p_o}}_{\text{GE effect}}\right] & \text{if route } h\text{'s origin} = o \text{ and } p_{\tilde{o}} = p_o \\ -\frac{\theta}{1-\theta}\left[\underbrace{\frac{0}{}}_{\text{PE effect}} + \underbrace{\frac{\partial g(K_h)}{\partial \tau_o}\frac{\tau_o}{g(K_h)} - \frac{1}{\theta}\frac{\partial p_{\hat{o}}}{\partial \tau_o}\frac{\tau_o}{p_{\hat{o}}}}_{\text{GE effect}}\right] & \text{if route } h\text{'s origin} \neq o \text{ and } p_{\tilde{o}} = p_{\hat{o}}. \end{cases}$$

PROOF: Retail output of goods sourced from route h with port fees, $\mathfrak{f}_o$, and tariffs, τ_o, is now given by

$$y_h = \left[\frac{\sigma\theta\chi_h p_o X_h}{(1+\tau_o)g(K_h)\left(\frac{\overline{p}_n}{1/\zeta}\right)^{1/\zeta}\left(\frac{\overline{p}_o+\mathfrak{f}_o}{1/\nu}\right)^{1/\nu}}\right]^{\theta/(1-\theta)} \tag{B.4}$$

and importers' sales are

$$p_o y_h = \left[\frac{\sigma\theta\chi_h p_o^{1/\theta} X_h}{(1+\tau_o)g(K_h)\left(\frac{\overline{p}_n}{1/\zeta}\right)^{1/\zeta}\left(\frac{\overline{p}_o+\mathfrak{f}_o}{1/\nu}\right)^{1/\nu}}\right]^{\theta/(1-\theta)}. \tag{B.5}$$

Differentiating both with respect to τ_o yields the elasticity of retail output and sales to tariffs. *Q.E.D.*

Proposition B.4 shows that an increase in tariffs does not necessarily reduce retail output. The impact depends on how it affects retail prices and congestion.

On one hand, an increase in tariffs on region o puts upward pressure on the retail price of goods from that same region, i.e., $\frac{\partial p_o}{\partial \tau_o}\frac{\tau_o}{p_o} \geq 0$. An increase in retail prices encourages importers to supply more foreign goods, but reduces consumer demand. On the other hand, an increase in tariffs decreases congestion along routes originating in the region where tariffs are levied, since $\frac{\partial g(K_h)}{\partial \tau_o}\frac{\tau_o}{g(K_h)} \leq 0$, and potentially leads to an increase in congestion on routes originating in other regions. A fall in congestion works like a productivity boost for shipping firms, allowing them to supply more goods in a shorter amount of time and at cheaper prices.

The cheaper shipping services in turn encourage importers to purchase more foreign goods since retail demand slopes down in the price of shipping services, $\frac{\partial y_h}{\partial p_h(j)}\frac{p_h(j)}{y_h} \leq 0$, which then leads to an increase in retail output.

PROPOSITION B.5: (IMPACT OF TARIFFS ON RETAIL SOURCING) The elasticity of importers' sourcing decision to tariffs is given by

$$\frac{\partial o_h}{\partial \tau_o}\frac{\tau_o}{o_h} = \lambda(1-\theta)\tau_o\left(\frac{\partial\left(p_{\tilde{o}}y_h\right)}{\partial \tau_o} - \sum_{h'} o_{h'}\frac{\partial\left(p_{\tilde{o}}y_{h'}\right)}{\partial \tau_o}\right),$$

where $p_{\tilde{o}}$ denotes the retail price of the goods sourced from the respective route.

PROOF: Importers source goods according to the discrete choice rule (4.16). Differentiating with respect to τ_o yields

$$\frac{\partial o_h}{\partial \tau_o}\frac{\tau_o}{o_h} = \lambda\tau_o\left(\frac{\partial \Pi_h}{\partial \tau_o} - \sum_{h'\in\mathcal{R}} o_{h'}\frac{\partial \Pi_{h'}}{\partial \tau_o}\right).$$

Retail profit from sourcing goods from route h is

$$\Pi_h = (1-\theta)\left[\frac{\sigma\theta\chi_h p_o^{1/\theta}X_h}{(1+\tau_o)g(K_h)\left(\frac{\overline{p}_n}{1/\zeta}\right)^{1/\zeta}\left(\frac{\overline{p}_o+\mathfrak{f}_o}{1/\nu}\right)^{1/\nu}}\right]^{\theta/(1-\theta)} - f. \tag{B.6}$$

As $\frac{\partial \Pi_h}{\partial \tau_o} = (1-\theta)\frac{\partial(p_o y_h)}{\partial \tau_o}$, substituting it into the expression above yields the result. *Q.E.D.*

The larger λ is, the more sensitive importers' sourcing decision is to tariff-induced differences in route profitability and the more imports reallocate away from the taxed routes toward routes whose profitability rises.

B.5. *The impact of port fees on real aggregate output*

PROPOSITION B.6: (IMPACT OF PORT FEES ON REAL AGGREGATE OUTPUT) The elasticity of real aggregate output Y to port fees $\mathfrak{f}_o$ on goods from region $o \in \mathcal{O}$ is given by

$$\begin{aligned}\frac{\partial Y}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{Y} = &\sum_{h\in\mathcal{R}}\left(\tfrac{\text{Route } h\text{'s share in}}{\text{aggregate output}}\right)\times\left(\frac{\partial y_h}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{y_h} + \frac{\partial o_h}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{o_h}\right)\\ &+\left(\tfrac{\text{Import share in}}{\text{aggregate output}}\right)\times\left(\frac{\partial \mathcal{I}}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{\mathcal{I}}\right) + \left(\tfrac{\text{Domestic share in}}{\text{aggregate output}}\right)\times\left(\frac{\partial d}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{d}\right),\end{aligned}$$

where $\frac{\text{Route } h\text{'s share in}}{\text{aggregate output}} = p_{o(h)}o_h y_h\mathcal{I}/(w+\Lambda+T)$ (with $o(h)$ route h's origin region), $\frac{\text{Domestic share in}}{\text{aggregate output}} = p_d d/(w+\Lambda+T)$, and $\frac{\text{Import share in}}{\text{aggregate output}} = 1 - \frac{\text{Domestic share in}}{\text{aggregate output}}$.

PROOF: The result follows from the same steps as the proof of Proposition 4.1. *Q.E.D.*

Propositions B.7 and B.8 below decompose the route-specific elasticities $\frac{\partial y_h}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{y_h}$ and $\frac{\partial o_h}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{o_h}$.

PROPOSITION B.7: (IMPACT OF PORT FEES ON IMPORTERS' OUTPUT AND SALES) The elasticity of retail output to port fees is

$$\frac{\partial y_h}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{y_h} = \begin{cases} -\frac{\theta}{1-\theta}\left[\underbrace{\frac{1}{\nu}\frac{\mathfrak{f}_o}{(\overline{p}_o+\mathfrak{f}_o)}}_{\text{PE effect}} + \underbrace{\frac{\partial g(K_h)}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{g(K_h)} - \frac{\partial p_o}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{p_o}}_{\text{GE effect}}\right] & \text{if route } h\text{'s origin} = o \\ -\frac{\theta}{1-\theta}\left[\underbrace{0}_{\text{PE effect}} + \underbrace{\frac{\partial g(K_h)}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{g(K_h)} - \frac{\partial p_{\tilde{o}}}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{p_{\tilde{o}}}}_{\text{GE effect}}\right] & \text{if route } h\text{'s origin} \neq o, \end{cases}$$

where $p_{\tilde{o}}$ is the retail price of the goods associated with route h when the region of origin is not o. The elasticity of retail sales to port fees is

$$\frac{\partial p_{\tilde{o}} y_h}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{p_{\tilde{o}} y_h} = \begin{cases} -\frac{\theta}{1-\theta}\left[\underbrace{\frac{1}{\nu}\frac{\mathfrak{f}_o}{(\overline{p}_o+\mathfrak{f}_o)}}_{\text{PE effect}} + \underbrace{\frac{\partial g(K_h)}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{g(K_h)} - \frac{1}{\theta}\frac{\partial p_o}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{p_o}}_{\text{GE effect}}\right] & \text{if route } h\text{'s origin} = o \text{ and } p_{\tilde{o}} = p_o \\ -\frac{\theta}{1-\theta}\left[\underbrace{0}_{\text{PE effect}} + \underbrace{\frac{\partial g(K_h)}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{g(K_h)} - \frac{1}{\theta}\frac{\partial p_{\tilde{o}}}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{p_{\tilde{o}}}}_{\text{GE effect}}\right] & \text{if route } h\text{'s origin} \neq o \text{ and } p_{\tilde{o}} = p_{\tilde{o}}. \end{cases}$$

PROOF: Differentiating Equations (B.4) and (B.5) with respect to $\mathfrak{f}_o$ yields the elasticity of retail output and sales to port fees. *Q.E.D.*

As was the case with tariffs, Proposition B.7 shows that an increase in port fees does not necessarily reduce retail output. First, the direct effect now operates through the input price $\mathfrak{f}_o/(\overline{p}_o+\mathfrak{f}_o)$ rather than the tariff wedge $\tau_o/(1+\tau_o)$, and is rescaled by the cost exponent $1/\nu$ because port fees are levied on shipping firms rather than retailers. This makes their direct transmission weaker than that of tariffs. Second, an increase in port fees on a particular region decreases congestion along routes originating in that same region and potentially raises congestion on routes originating in other regions. It can also lead to an increase in retail prices as port fees increase the cost of shipping services.

PROPOSITION B.8: (IMPACT OF PORT FEES ON RETAIL SOURCING) The elasticity of importers' sourcing decision to port fees is given by

$$\frac{\partial o_h}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{o_h} = \lambda(1-\theta)\mathfrak{f}_o\left(\frac{\partial\left(p_{\tilde{o}} y_h\right)}{\partial \mathfrak{f}_o} - \sum_{h'\in\mathcal{R}} o_{h'}\frac{\partial\left(p_{\tilde{o}} y_{h'}\right)}{\partial \mathfrak{f}_o}\right),$$

where $p_{\tilde{o}}$ denotes the retail price of the goods sourced from the respective route.

PROOF: Importers source goods according to the discrete choice rule (4.16). Differentiating with respect to $\mathfrak{f}_o$ yields

$$\frac{\partial o_h}{\partial \mathfrak{f}_o}\frac{\mathfrak{f}_o}{o_h} = \lambda \mathfrak{f}_o \left(\frac{\partial \Pi_h}{\partial \mathfrak{f}_o} - \sum_{h' \in \mathcal{R}} o_{h'} \frac{\partial \Pi_{h'}}{\partial \mathfrak{f}_o} \right).$$

Retail profit from sourcing goods from route h is given by Equation (B.6). As $\frac{\partial \Pi_h}{\partial \mathfrak{f}_o} = (1-\theta)\frac{\partial (p_o y_h)}{\partial \mathfrak{f}_o}$, substituting it into the expression above yields the result. *Q.E.D.*

The larger λ is, the more sensitive importers' sourcing decisions are to fee-induced differences in route profitability and the more imports reallocate away from the fee-affected routes toward routes whose profitability rises.

C: QUANTITATIVE ANALYSIS

This appendix reports the estimation results behind the calibration of Section 5, the baseline simulation, the validation against spot freight rates, and additional results for the counterfactual experiments of Sections 6–9.

C.1. *Calibration and baseline simulation*

TABLE A1

AVERAGE NUMBER OF TRIPS AND AVERAGE SPEED

	Log(Average Number of Trips per Vessel)			
	OLS		WLS	
	(1)	(2)	(3)	(4)
Log(Average Sailing Speed)	-0.197	0.035	0.674***	0.663**
	(0.140)	(0.148)	(0.233)	(0.327)
Route FE	Yes	Yes	Yes	Yes
Year & Month FEs	No	Yes	No	Yes
Weights	None	None	DWT	DWT
Observations	666	666	666	666
R^2	0.901	0.955	0.853	0.940

Note: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

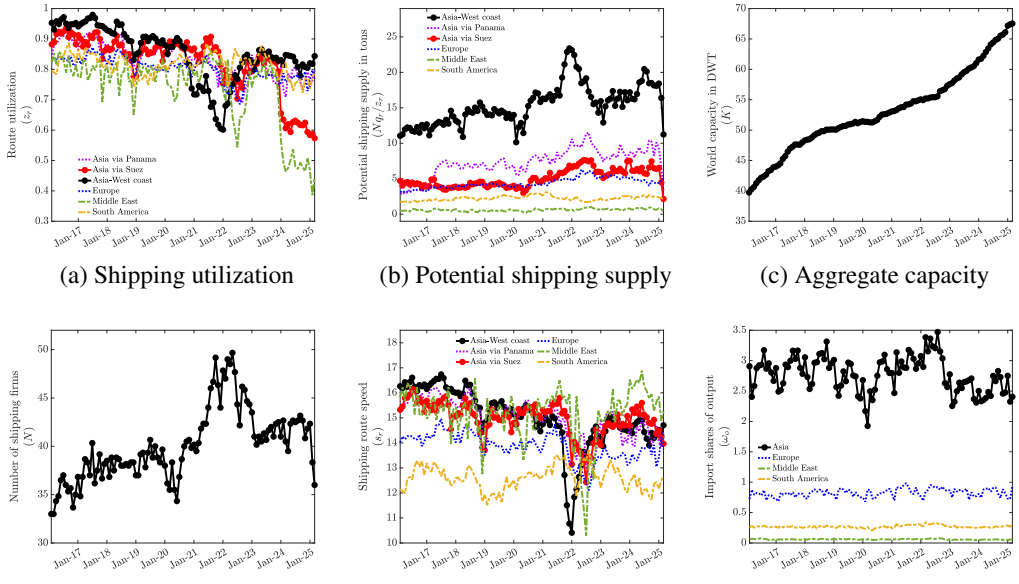
TABLE A2

TOTAL PORT CONGESTION TIME AND TOTAL NUMBER OF PORT CALLS

	Log(Total Congestion Time)					
	OLS		WLS			
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Total Port Calls)	1.409***	1.750***	0.868***	1.296***	1.338***	1.437***
	(0.074)	(0.076)	(0.080)	(0.085)	(0.084)	(0.088)
Port FE	Yes	Yes	Yes	Yes	Yes	Yes
Year & Month FEs	No	Yes	No	Yes	No	No
Year × Month FEs	No	No	No	No	Yes	Yes
Weights	None	None	Port Calls	Port Calls	Port Calls	Port Calls
Observations	5,139	5,139	5,139	5,139	5,139	3,471
R^2	0.751	0.776	0.760	0.784	0.789	0.812

Note: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Column (6) excludes 2020–2022.

FIGURE A1.—Targeted time series, February 2016–March 2025

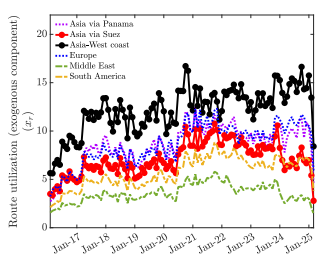


(a) Shipping utilization (b) Potential shipping supply (c) Aggregate capacity

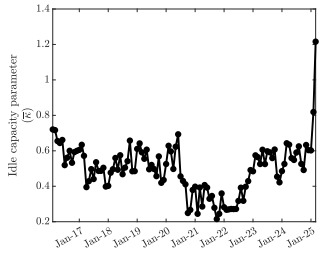
(d) Number of shipping firms (e) Average sailing speeds (f) Import shares

Note: Panel A1a shows our measure of shipping capacity utilization defined as the ratio of effective shipping supply to potential shipping supply for the six shipping routes. Panel A1b shows the potential shipping supply in tons for the six shipping routes. Panel A1c shows the world capacity in DWT. Panel A1d shows the number of shipping firms. Panel A1e shows the average sailing speeds along the six shipping routes. Panel A1f shows the import shares of output from Asia, Europe, the Middle East, and South America. The targeted measure of importers is equal to one over the same period. All these model-implied time series match the data.

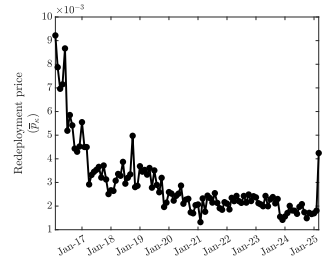
FIGURE A2.—Evolution of the underlying exogenous variables, February 2016–March 2025



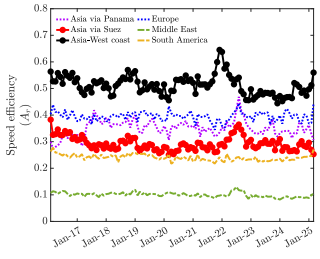
(a) Route utilization (exog)



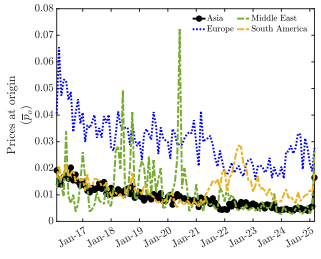
(b) Idle capacity parameter



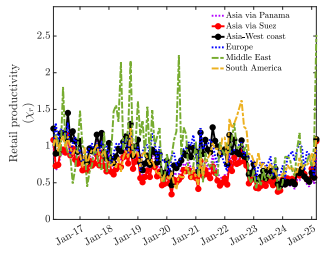
(c) Redeployment price



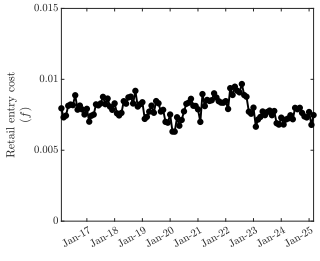
(d) Speed efficiency



(e) Origin prices



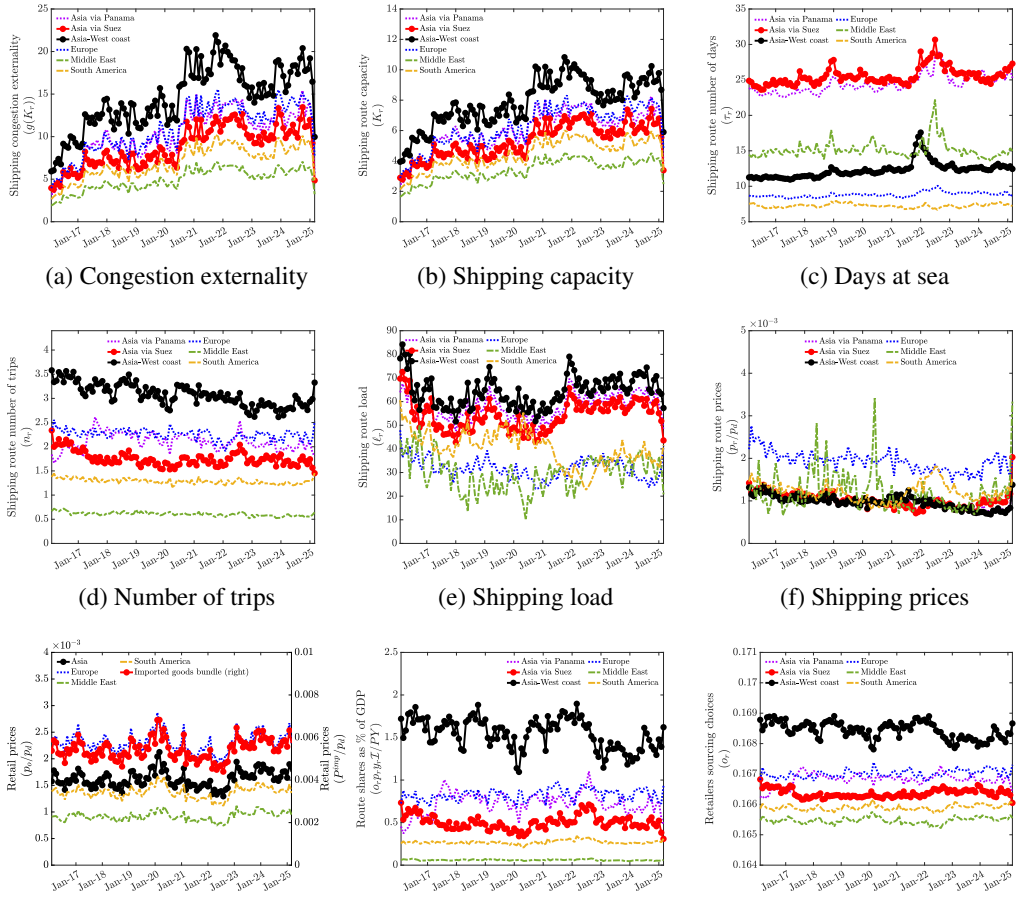
(f) Retail productivity



(g) Retail entry cost

Note: Panel A2a shows the exogenous component of the shipping capacity utilization. Panel A2b shows the idle capacity parameter. Panel A2c shows the price of redeployment. Panel A2d shows the speed efficiency. Panel A2e shows the origin prices of goods produced in Asia, Europe, the Middle East, and South America. Panel A2f shows importers' productivity. Panel A2g shows the entry cost in the retail market. These exogenous variables match the time series displayed in Figure A1 together with the constraint on the measure of importers.

FIGURE A3.—Evolution of other endogenous variables, February 2016–March 2025



(a) Congestion externality (b) Shipping capacity (c) Days at sea

(d) Number of trips (e) Shipping load (f) Shipping prices

(g) Retail prices (h) Route shares as % of GDP (i) Sourcing probabilities

Note: Panel A3a shows the congestion externality in shipping productivity. Panel A3b shows the aggregate shipping capacity. Panel A3c shows the average number of days spent on voyages to the U.S. Panel A3d shows the average number of trips. Panel A3e shows the shipping load. Panel A3f shows the price of shipping services. Panel A3g shows the retail price of imported goods. Panel A3h shows the route shares as a fraction of GDP. Panel A3i shows importers' sourcing probabilities.

C.2. *Shipping utilization and spot container freight rates*

Equations (4.8) and (4.10) imply that the marginal cost of shipping is inversely proportional to utilization, so that, holding the markup fixed, the log price of shipping services rises one-for-one with shipping pressure, $\psi_{rt} \equiv -\log z_{rt}$. This appendix complements the descriptive evidence in Section 5.3 with regressions of log spot rates on shipping pressure at the route-month level for the two disruption episodes studied in Section 6.

We use monthly Shanghai Containerized Freight Index (SCFI) spot rates from Clarksons' Shipping Intelligence Network (Clarksons Research Services Limited, 2025). For the Pandemic episode, we use the Shanghai–U.S. West Coast series (series 98514) from January 2020 through December 2022 and match it to APAC–USWC utilization. For the Red Sea episode, we use the Shanghai–U.S. East Coast series (series 98515) from August 2023 through March 2025 and match it to APAC–USEC utilization for the Suez/Cape route. Both SCFI series are measured in U.S. dollars per forty-foot equivalent unit (FEU). Monthly utilization, z_{rt}, is the arithmetic mean of the paper's weekly route-level measure within each calendar month.

For each episode, we estimate

$$\log(\mathrm{SCFI}_{rt}) = b_0 + b_1 \psi_{rt} + b_2\, t + \eta_{rt}, \tag{C.1}$$

where t is a linear monthly trend. Standard errors use a Bartlett-kernel Newey–West covariance estimator with three monthly lags and the $n/(n-k)$ small-sample correction.

Table A3 reports $b_1 = 3.460$ (standard error 0.961, $p = 0.0003$) for the pandemic episode and $b_1 = 3.634$ (standard error 0.889, $p < 0.0001$) for the Red Sea episode. Conditional on the trend, a 1% decline in shipping utilization is associated with a roughly 3.5% increase in the route's SCFI spot rate. The elasticities are positive and statistically significant in both episodes, even though the episodes differ in origin, consistent with the model's marginal-cost channel. The estimated elasticities exceed the one-for-one benchmark: as discussed in Section 5.3, spot rates reflect only the freight component, whereas $p_{r,t}$ is a delivered price inclusive of the goods. Hence, we interpret the comovement as validating the sign and timing of the model's price response rather than its level. The estimates are descriptive and not used in the calibration: SCFI measures spot rates per FEU, and its East Coast series aggregates services beyond the Suez route.

TABLE A3

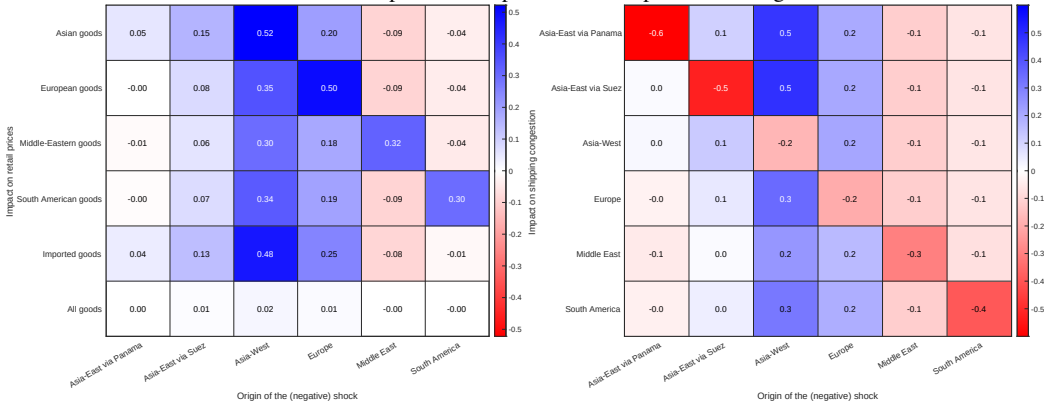
SHIPPING UTILIZATION AND SPOT CONTAINER FREIGHT RATES

	Log(SCFI Spot Freight Rate)	
	Pandemic (1)	Red Sea (2)
Shipping pressure	3.460***	3.634***
	(0.961)	(0.889)
p-value	0.0003	< 0.0001
Sample window	Jan. 2020–Dec. 2022	Aug. 2023–Mar. 2025
Linear monthly trend	Yes	Yes
HAC bandwidth	3 months	3 months
Observations	36	20

Note: The dependent variable is the log SCFI spot rate in dollars per FEU.
Monthly utilization is the arithmetic mean of the paper's weekly measure; shipping pressure is minus its log.
Columns (1) and (2) match SCFI rates to APAC–USWC and APAC–USEC utilization, respectively.
Newey–West standard errors use a Bartlett kernel, three lags, and the $n/(n-k)$ correction.
The estimates are descriptive. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

C.3. *Additional results for Proposition 4.1*

FIGURE A4.—Impact of disruptions on retail prices and congestion

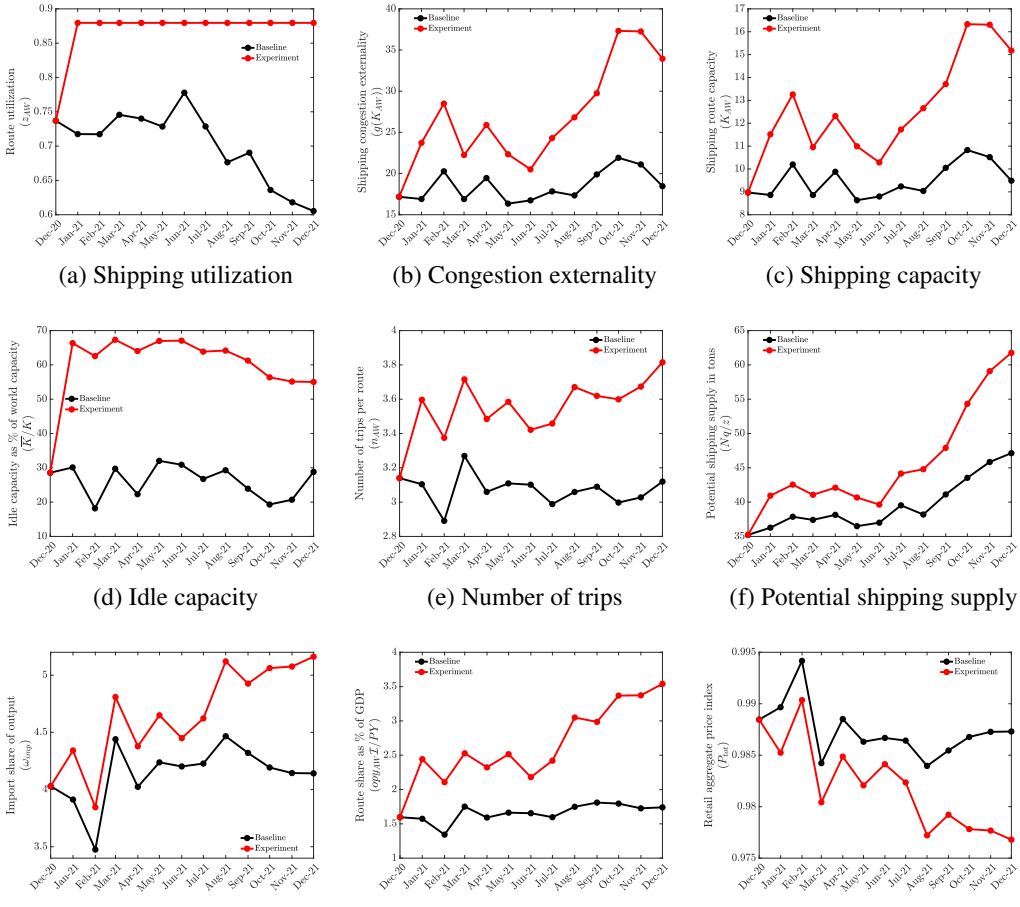


(a) Impact on retail prices (b) Impact on congestion

Note: Elasticities of retail prices (a) and congestion (b) to route-r disruptions, $-[\partial \log(p_o)/\partial \log(X_r)]$ and $-[\partial \log(g(K_h))/\partial \log(X_r)]$.

C.4. *Impact of the 2021 West Coast port crisis (Section 6.2)*

FIGURE A5.—Evolution of selected aggregates: baseline vs. low shipping disruptions along the Asia to the West Coast route, Dec 2020–Dec 2021



(a) Shipping utilization (b) Congestion externality (c) Shipping capacity

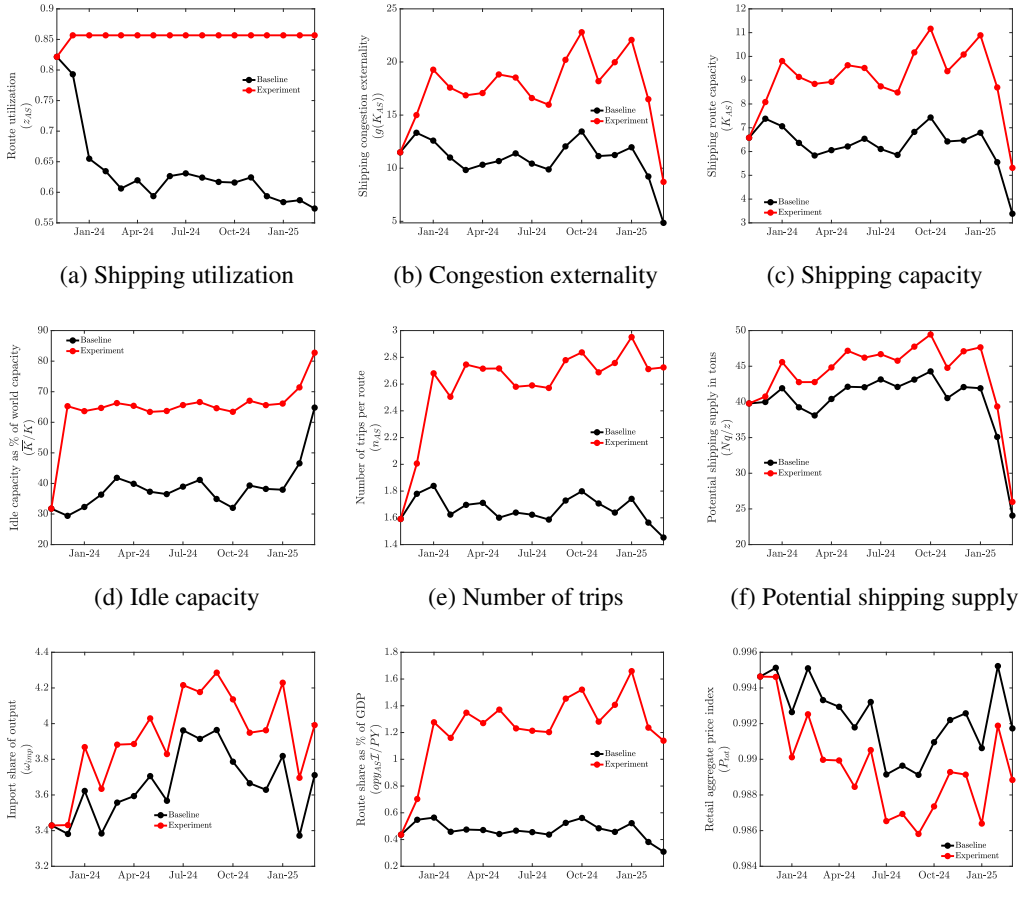
(d) Idle capacity (e) Number of trips (f) Potential shipping supply

(g) Import share (h) Route share (i) Aggregate price index

Note: Panel (a) shows the shipping utilization along the Asia to West Coast route. Panel (b) shows the congestion externality along the Asia to West Coast route. Panel (c) shows the shipping capacity allocated to the Asia to West Coast route. Panel (d) shows the idle capacity as a share of world capacity. Panel (e) shows the number of trips realized along the Asia to West Coast route. Panel (f) shows the U.S.-bound potential shipping supply. Panel (g) shows the import share of GDP. Panel (h) shows the Asia to West Coast route share of GDP. Panel (i) shows the aggregate price index.

C.5. *Impact of the 2024 security threats in the Red Sea (Section 6.3)*

FIGURE A6.—Evolution of selected aggregates: baseline vs. low shipping disruptions along the Asia to the East Coast route via the Suez Canal, December 2023–March 2025



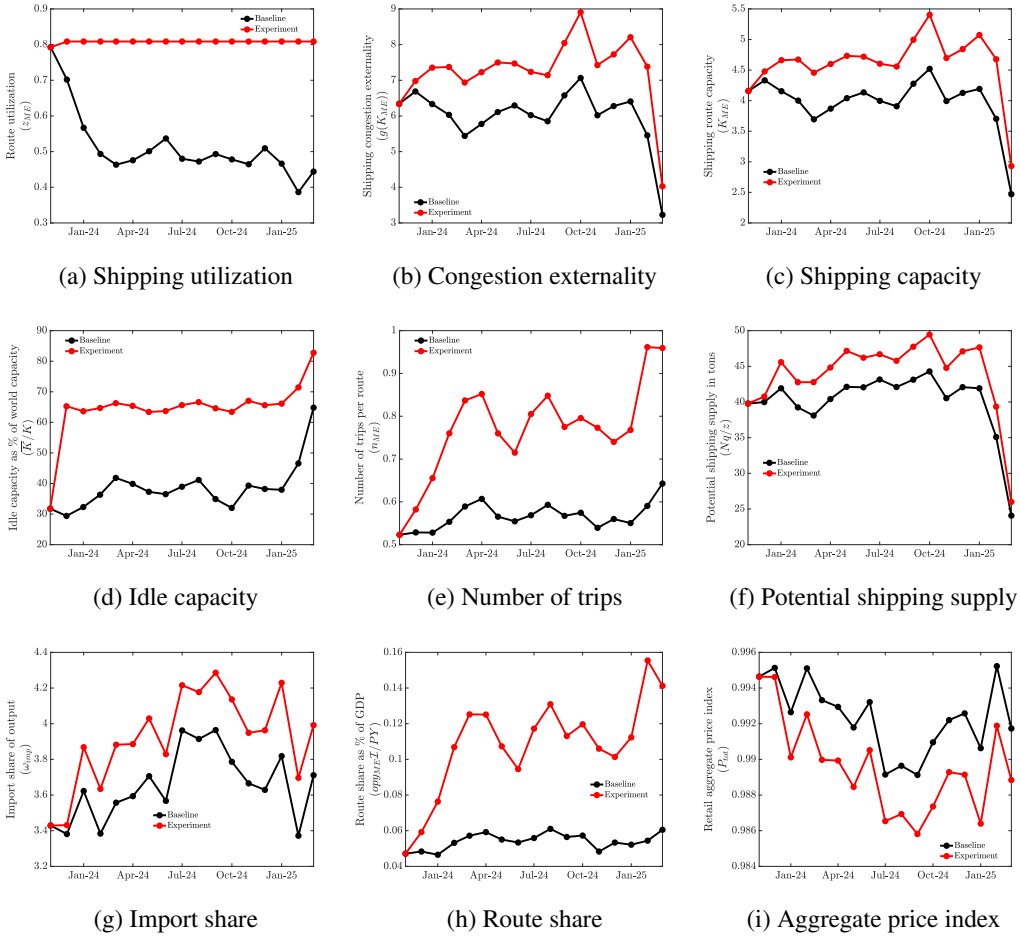
(a) Shipping utilization (b) Congestion externality (c) Shipping capacity

(d) Idle capacity (e) Number of trips (f) Potential shipping supply

(g) Import share (h) Route share (i) Aggregate price index

Note: Panel (a) shows the shipping utilization along the Asia to East Coast via the Suez Canal route. Panel (b) shows the congestion externality along the Asia to East Coast via the Suez Canal route. Panel (c) shows the shipping capacity allocated to the Asia to East Coast via the Suez Canal route. Panel (d) shows the idle capacity as a share of world capacity. Panel (e) shows the number of trips realized along the Asia to East Coast via the Suez Canal route. Panel (f) shows the U.S.-bound potential shipping supply. Panel (g) shows the import share of GDP. Panel (h) shows the Asia to East Coast via the Suez Canal route share of GDP. Panel (i) shows the aggregate price index.

FIGURE A7.—Evolution of selected aggregates: baseline vs. low shipping disruptions along the Middle East to the East Coast route, December 2023–March 2025



(a) Shipping utilization (b) Congestion externality (c) Shipping capacity

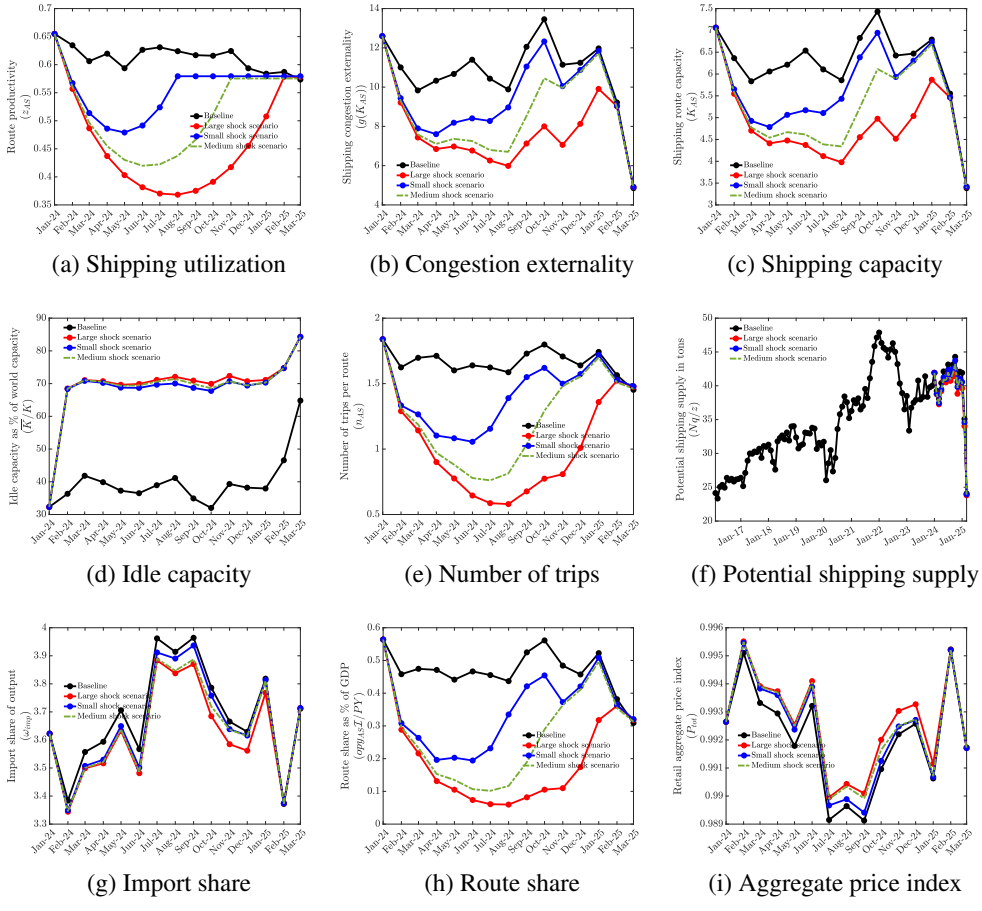
(d) Idle capacity (e) Number of trips (f) Potential shipping supply

(g) Import share (h) Route share (i) Aggregate price index

Note: Panel (a) shows the shipping utilization along the Middle East to the East Coast route. Panel (b) shows the congestion externality along the Middle East to the East Coast route. Panel (c) shows the shipping capacity allocated to the Middle East to the East Coast route. Panel (d) shows the idle capacity as a share of world capacity. Panel (e) shows the number of trips realized along the Middle East to the East Coast route. Panel (f) shows the U.S.-bound potential shipping supply. Panel (g) shows the import share of GDP. Panel (h) shows the Middle East to the East Coast route share of GDP. Panel (i) shows the aggregate price index.

C.6. *Impact of the 2024 Navy intervention in the Red Sea (Section 7)*

FIGURE A8.—Evolution of selected aggregates: baseline vs. no Navy support along the Asia to the East Coast route via the Suez Canal, January 2024–March 2025



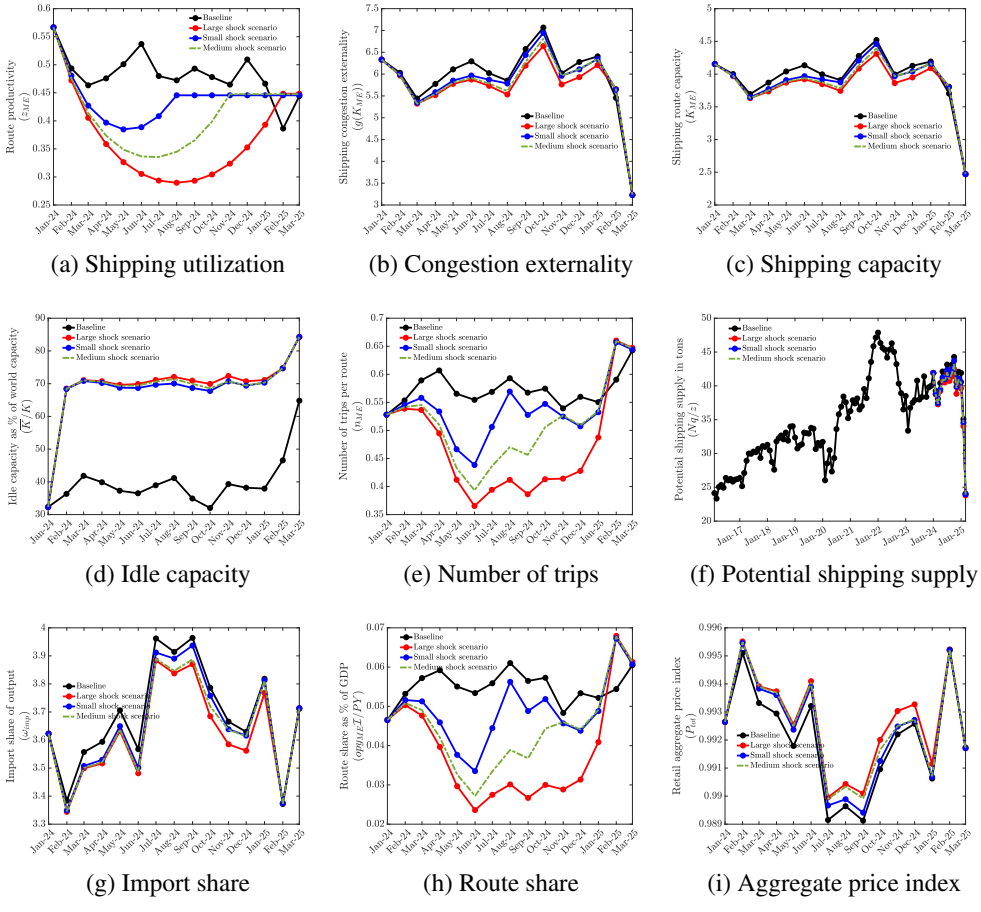
(a) Shipping utilization (b) Congestion externality (c) Shipping capacity

(d) Idle capacity (e) Number of trips (f) Potential shipping supply

(g) Import share (h) Route share (i) Aggregate price index

Note: Panel (a) shows the shipping utilization along the Asia to East Coast via the Suez Canal route. Panel (b) shows the congestion externality along the Asia to East Coast via the Suez Canal route. Panel (c) shows the shipping capacity allocated to the Asia to East Coast via the Suez Canal route. Panel (d) shows the idle capacity as a share of world capacity. Panel (e) shows the number of trips realized along the Asia to East Coast via the Suez Canal route. Panel (f) shows the U.S.-bound potential shipping supply. Panel (g) shows the import share of GDP. Panel (h) shows the Asia to East Coast via the Suez Canal route share of GDP. Panel (i) shows the aggregate price index.

FIGURE A9.—Evolution of selected aggregates: baseline vs. no Navy support along the Middle East to the East Coast route, January 2024–March 2025



(a) Shipping utilization (b) Congestion externality (c) Shipping capacity

(d) Idle capacity (e) Number of trips (f) Potential shipping supply

(g) Import share (h) Route share (i) Aggregate price index

Note: Panel (a) shows the shipping utilization along the Middle East to the East Coast route. Panel (b) shows the congestion externality along the Middle East to the East Coast route. Panel (c) shows the shipping capacity allocated to the Middle East to the East Coast route. Panel (d) shows the idle capacity as a share of world capacity. Panel (e) shows the number of trips realized along the Middle East to the East Coast route. Panel (f) shows the U.S.-bound potential shipping supply. Panel (g) shows the import share of GDP. Panel (h) shows the Middle East to the East Coast route share of GDP. Panel (i) shows the aggregate price index.

C.7. *Robustness of the welfare estimates*

Table A4 reports the welfare estimates when each structural parameter is perturbed by $\pm 10\%$ around its baseline value, recalibrating the internally determined parameters in every row so that the baseline economy continues to reproduce the same targets and time paths. The welfare cost falls with σ, rises with θ, ϱ, and ρ, and is unaffected by the remaining parameters.

Panel A varies the shipping side. A larger shipping markup (i.e., a lower σ) amplifies the pass-through of the same efficiency loss into shipping and import prices, and the amplified price swing propagates through entry, capacity, and consumption. The welfare cost of disruptions therefore rises with shipping market power. The estimates are instead invariant to the

TABLE A4

ROBUSTNESS OF THE WELFARE ESTIMATES

Specification	Equivalent variation (% of output)		
	No 2021 port disruptions	No 2024 security threats	No 2024 Navy intervention
Baseline	0.693	0.354	-0.058
Panel A: Shipping markup and costs			
Lower shipping markup ($\sigma = 0.88$)	0.592	0.322	-0.059
Higher shipping markup ($\sigma = 0.72$)	0.854	0.405	-0.055
Lower trip cost curv. ($\zeta = 1.8$, $\nu = 2.2$)	0.693	0.354	-0.058
Higher trip cost curv. ($\zeta = 2.2$, $\nu = 1.8$)	0.693	0.354	-0.058
Panel B: Importers' markup and sourcing			
Lower retail markup ($\theta = 0.88$)	0.881	0.586	-0.030
Higher retail markup ($\theta = 0.72$)	0.582	0.271	-0.071
Lower sourcing variance ($\lambda = 0.9$)	0.692	0.354	-0.058
Higher sourcing variance ($\lambda = 1.1$)	0.693	0.354	-0.058
Panel C: Consumer demand			
Lower sub. dom. vs. imp. goods ($\varrho = 0.45$)	0.634	0.327	-0.054
Higher sub. dom. vs. imp. goods ($\varrho = 0.55$)	0.774	0.393	-0.063
Lower sub. among imp. goods ($\rho = 0.675$)	0.684	0.351	-0.059
Higher sub. among imp. goods ($\rho = 0.825$)	0.698	0.357	-0.057
Panel D: Pre-disruption anchor date			
Alternative anchor date (one month earlier)	0.758	0.309	-0.058

Notes: The table reports the consumption equivalent variation (as a percentage of output) for the counterfactual economies without the 2021 West Coast port disruptions, the 2024 Red Sea security threats, and the 2024 U.S. Navy intervention (central estimate), under alternative parameter values. Each row recalibrates the internally determined parameters and simulates the model so that the baseline economy continues to reproduce the same data targets and time paths. Panels A–C perturb each parameter by $\pm 10\%$ around its baseline value. For the Navy intervention scenario, the higher substitution parameter among imported goods, ρ, is perturbed by 5% only for numerical stability over the evaluation window. In the last column a negative entry means consumers are worse off in the economy without the intervention, so its magnitude is the value of the Navy's protection.

parameters that only shape the supply-side decomposition of measured efficiency—the trip-cost curvature (ζ), the trip technology (α), and the congestion function (ϕ, ε_k)—because the counterfactuals feed the measured efficiency paths z_r directly into the simulations.

Panel B varies the importers' problem. The returns to scale θ moves the estimates in the opposite direction to σ even though both parameters govern a markup. A higher θ brings the retail sector closer to constant returns and thins importers' margins, so trade, retail entry, and route capacity all become more responsive to shipping conditions, and the same disruption suppresses more activity. A lower θ makes import supply inelastic, and part of the shock is absorbed in inframarginal profits rather than in quantities. The sourcing dispersion λ barely moves the estimates. It reallocates importers across routes, but its scope is limited.

Panel C varies consumer demand. Intuitively, more substitutable goods should make disruptions less costly, as households can more easily switch away from the goods that become expensive. Here the opposite happens, because switching is itself distorted. With a higher ϱ, households respond to more expensive imports by buying domestic goods instead, but as imports sell above marginal cost, every unit given up was worth more than it cost to supply and the switch destroys that surplus. With a higher ρ, households switch toward origins served by less-affected routes, which dampens the price increase of the affected goods but congests the remaining routes. So the disruption spreads to all import prices instead of staying contained.

Finally, the alternative anchor in Panel D changes the implied depth of each episode, raising the 2021 estimate and lowering the 2024 one.

C.8. *Impact of port infrastructure improvement*

Our final appendix analyzes the general equilibrium effect of port infrastructure improvements that alleviate congestion. We model the improvement by lowering the value of ε_k, the congestion shifter to 0.9. Because ε_k scales the congestion-time relationship estimated in Section 5, the experiment has a direct empirical interpretation: holding the number of port calls fixed, a lower ε_k reduces the hours container ships wait in a port's anchorage before berthing by 10%. A 10% improvement is modest compared with the congestion that built up during the 2021–2022 supply chain crisis, when aggregate shipping utilization fell from roughly 89% to below 70%. Table A5 contrasts the counterfactual economy with the baseline economy matched to March 2025 aggregates.

TABLE A5

THE IMPACT OF PORT INFRASTRUCTURE IMPROVEMENTS

		Baseline (1)	$\downarrow \varepsilon_k$ (2)	$\Delta\%$ (3)
Shipping market				
Q_r	Effective shipping supply (in DWT), total	18.942	20.767	9.6
$g(K_r)$	Congestion factor, *Asia-East Coast via Panama*	5.701	5.439	-4.6
	Asia-East Coast via Suez	4.855	4.632	-4.6
	Asia-West Coast	9.967	9.510	-4.6
	Europe to East Coast	7.635	7.286	-4.6
	Middle East to East Coast	3.228	3.080	-4.6
	South America to East Coast	4.864	4.641	-4.6
p_r	Shipping prices, *Asia-East Coast via Panama* ($\times 100$)	0.146	0.139	-4.6
	Asia-East Coast via Suez ($\times 100$)	0.203	0.193	-4.6
	Asia-West Coast ($\times 100$)	0.138	0.132	-4.6
	Europe to East Coast ($\times 100$)	0.196	0.187	-4.6
	Middle East to East Coast ($\times 100$)	0.333	0.318	-4.6
	South America to East Coast ($\times 100$)	0.155	0.148	-4.6
N	Number of shipping firms	36.000	37.661	4.6
Retail market				
$\mathcal{I}$	Measure of retailers	1.000	1.046	4.6
ω_o	Import share, total	0.037	0.039	4.6
P^{imp}	Retail price of imported goods ($\times 100$)	0.574	0.547	-4.5
P^{tot}	Retail price of all goods	0.992	0.990	-0.2
ε_{ev}	Welfare: equivalent variation (% of output)		0.211	

Note: The experiment is based on the model calibrated to March 2025 data. Route-level responses are nearly uniform as ε_k affects all routes uniformly.

The improvement raises the productivity of shipping firms, allowing them to deliver more goods faster. The effective shipping supply increases by about 9.6% across all routes. Part of the gain is competed away. Firms reallocate ships toward decongested routes, so although ε_k falls by 10%, the congestion factor falls by only 4.6%.

As shipping productivity rises, the prices of shipping services and of imported goods fall by 4.6 and 4.5%, and the import share increases by 4.6%. Consumers would need an additional 0.21% of output in the baseline economy to be as well off as in the economy with better port infrastructure. The Infrastructure Investment and Jobs Act of 2021 allocated about $17 billion (about 0.07% of GDP) spread over several years to ports and waterways (The White House, 2021, American Society of Civil Engineers, 2025), which suggests that port decongestion is a high-return investment.