November 2026

Maritime Vulnerabilities and Geopolitical Disruptions in Global Shipping[1]

Simon Fuchs
Federal Reserve Bank of Atlanta
Monash University
CESifo

Fernando Leibovici
Federal Reserve Bank of Dallas
Southern Methodist University

Woan Foong Wong
University of Oregon
CEPR & NBER
CESifo

Abstract

We study how dependence on maritime routes and foreign vessel capacity shapes vulnerability to disruptions. We develop an analytical framework linking trade exposure, shipping-cost changes, and producer-price adjustment to welfare losses and to the incentives to initiate disruptions. These incentives depend on the target's loss relative to the economic cost borne by the disruptor. Combining data on vessel registries, port rotations, and bilateral trade, we measure each country's exposure to shipping lanes, its scope for rerouting, and its reliance on foreign-operated and foreign-built vessels. We then quantify the welfare impact of maritime disruptions in a spatial equilibrium model with endogenous routing and vessel deployment, which can account for the observed rerouting and cross-market freight rate responses during the Red Sea crisis. We find that exposure need not imply vulnerability: the cost of disrupting a route depends on whether ships can sail around it. U.S. losses are spread across distant chokepoints, while China's are larger and concentrated in East Asia, and the largest losses fall on small economies near major chokepoints. Disruptions that affect several nearby chokepoints at once, as around Taiwan and the South China Sea, remove these alternatives, so a disruption there that restricts only some countries' cargo can impose large losses on them at little cost to others. Dependence on foreign vessels also shapes vulnerability: because routes compete for the same fleet, a withdrawal of Chinese-operated or Chinese-built vessels from routes serving the United States imposes larger losses on several Asian economies than on the United States itself.

KEYWORDS: Global shipping, foreign dependence, transport networks, trade costs, geoeconomics

[1]Contact: simon.fuchs@atl.frb.org, fleibovici@gmail.com, and wfwong@uoregon.edu.

We thank Kerem Coşar, Nathalie El-Bazzal, Keith Head, Myrto Kalouptsidi, and Theodore Papageorgiou for their insightful discussions, Dawn Chinagorom-Abiakalam and Valerie Nguyen for excellent research assistance, and participants at the 2026 NBER Summer Institute International Trade and Investment meeting, Kiel Institute World Trade Dialogue, TIGN, 2026 Yale Cowles Summer Conference, FRB Atlanta, 2025 NBER Transportation Networks Conference, SEA, Keio University, the World Bank, Asia School of Business, the University of Melbourne, Monash University, National University of Singapore, University of Oklahoma, the BSE Summer Forum Workshop on Geography, Trade, and Growth, and the 2nd International Conference of the Georgian Economic Association for helpful comments. The views expressed herein are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of Dallas, the Federal Reserve Bank of Atlanta, the Federal Reserve System, or the Board of Governors. All errors are our own.

1 Introduction

Dependence on foreign-controlled inputs has long been viewed as a source of geopolitical risk. Hirschman (1945) showed how concentrated trade relationships create asymmetric dependence and alter bargaining power. Recent work on geoeconomics examines how these relationships operate through production networks, sanctions, and cross-border infrastructure (Farrell and Newman, 2019, 2023; Clayton, Maggiori and Schreger, 2026; Mohr and Trebesch, 2025). Maritime transportation is central to this dependence. More than 80 percent of world trade by volume moves by sea (UNCTAD, 2024). Maritime geography and scale economies concentrate this traffic on a small set of routes, ports, and chokepoints. Moreover, shipbuilding is concentrated in a few countries, and a small number of large carriers operates much of the fleet. While this concentration lowers shipping costs in normal times, it can also expose trade to disruptions along key routes and to restrictions on the vessels that serve them.

Governments have long treated both maritime routes and vessel capacity as strategic. Competition over maritime geography is reflected in the number and concentration of militarized disputes around canals, straits, as well as heavily used sea lanes (Figures 1a and 1b, respectively). Efforts to retain national control extend to the vessels serving these routes: 91 countries have implemented cabotage laws reserving domestic shipping for vessels that meet national requirements on flag, ownership, construction, or crewing. Overall, these laws cover 80 percent of the world's coastlines (Seafarers' Rights International, 2018). This tension between national control and international dependence motivates our central question: when does reliance on foreign routes and vessels make countries economically vulnerable?

We study how dependence on maritime routes and foreign shipping capacity shapes the welfare consequences of disruptions and the incentives to initiate them. We develop an analytical framework linking trade exposure, shipping-cost responses, and equilibrium price adjustment to country welfare. We then use global containership, port rotation, and bilateral trade data to measure countries' exposure to shipping routes, the scope for rerouting, and dependence on foreign operators and builders. We bring these elements together in a spatial equilibrium model with endogenous routing and vessel deployment, allowing us to evaluate geographic disruptions, restrictions on shipping capacity, and policies intended to reduce maritime dependence.

Our results show that a country's exposure to a route or to a set of vessels does not, by itself, determine the welfare cost of losing access to them; this cost also depends on the alternatives available. First, across the global shipping network, the costliest disruptions occur at chokepoints with limited alternatives, such as Bab el-Mandeb, the approach to Gibraltar, and the Strait of Malacca, and the largest losses fall on small economies along these routes. Second, when several nearby chokepoints are disrupted together, losses depend on whether the affected cargo can be rerouted. Taiwan and the South China Sea provide a natural setting: major routes between East Asia and the rest of the world pass through a dense cluster of nearby passages, from the Taiwan Strait to the Straits of Malacca and Sunda, that can substitute for one another. There, a disruption limited to cargo not traded with China cuts off Taiwan's trade with other partners, imposing large losses on Taiwan and nearby economies while leaving China's welfare nearly

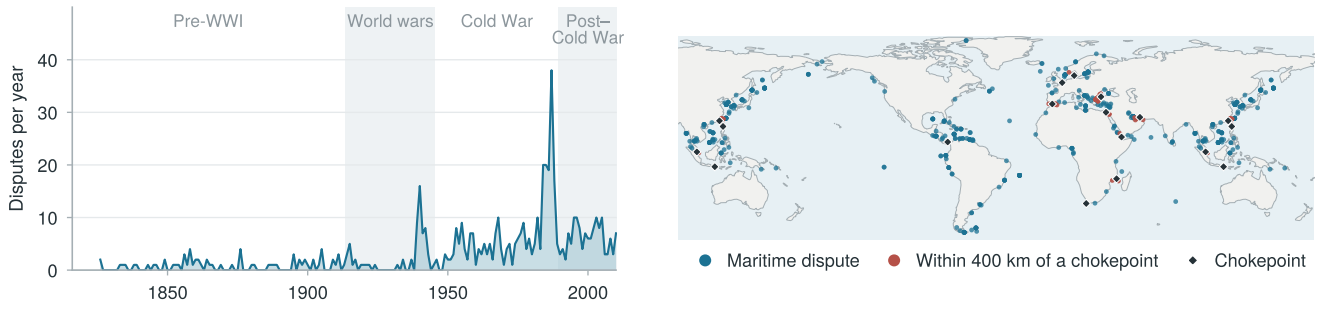
unchanged. Most China-related cargo can instead detour around any single passage, so restrictions at other chokepoints make such a disruption costly for China only when several are combined, and they shift costs to third countries. Third, the same holds for vessels. Restricting Chinese-operated or Chinese-built capacity on routes serving the United States imposes larger losses on several other economies than on the United States, as the remaining fleet is redeployed. Fees that shift U.S. trade away from these vessels reduce the additional loss from such a disruption, but by less than they cost in normal times.

We begin by developing a sufficient-statistics framework that connects maritime dependence to the welfare consequences of a disruption. Building on the approach in Allen, Fuchs and Wong (2025), which allows for endogenous transport costs, we extend their results to derive a sufficient statistic for the welfare costs of disruptions in a multi-country spatial Armington economy. To first order, a maritime disruption affects a country's welfare through two terms. The first is the change in shipping costs on each connection, weighted by the value of the country's purchases shipped through it. The second captures changes in producer prices, which affect both what households pay for goods and what they earn from production. The change in shipping costs is itself an equilibrium outcome: it depends on the disruption, the alternatives available to cargo and vessels, and how freight rates and vessel deployment adjust. A disruption to a heavily used route can therefore impose little loss if substitutes are readily available, while a disruption to a less-used route can be costly if alternatives are scarce. This distinction guides what we measure in the data and what we quantify in the model.

We then examine how these welfare consequences shape the incentives to initiate a disruption. A country may value imposing losses abroad but incur economic costs from doing so. These include its own welfare loss, the cost of implementing the restriction, and any losses in shipping earnings not already included in welfare. We distinguish the target's welfare loss from the disruptor's cost. The ratio of the two determines how attractive the disruption is, taking into account the economic cost to the disruptor. Dependence can therefore operate in both directions: foreign shipping capacity may expose a country's trade to restrictions, but the earnings and trade associated with supplying that capacity can also make those restrictions costly to impose. The framework thus connects the welfare consequences of maritime disruptions to the economic trade-off faced by a country considering whether to initiate one.

We use the framework to guide our measurement of maritime dependence. This requires tracing not only which countries trade with one another, but also the routes and vessels connecting them. Shipments may pass through several foreign ports before reaching their destination, creating dependence on connections that never touch a country's own shores. We combine bilateral trade records with vessel registries and scheduled port rotations to estimate these flows through the global containership network. For each lane, we measure countries' trade exposure, the additional shipping distance required to bypass local disruptions, and the shares of capacity associated with foreign operators and builders.

We find that countries' trade exposure extends well beyond their own ports. U.S. exposure is concentrated on trans-Pacific and North Atlantic routes, while Chinese exposure is concentrated around East Asian hubs and Asia–Europe routes. Both countries depend on foreign gateways and lanes between third



(a) Annual maritime militarized disputes, 1816–2010

(b) Geographic distribution of disputes

Figure 1. Maritime militarized disputes. Panel (a) reports annual counts of water-classified militarized interstate disputes. Panel (b) maps their locations; red markings show disputes within 400 km of major chokepoints. Source: MID 5.0 (Palmer et al., 2022) and MIDLOC 2.1 (Braithwaite, 2010; Bezerra and Braithwaite, 2019).

countries, so bilateral trade relationships do not fully describe their reliance on maritime transportation. We also document extensive exposure to routes through foreign waters, including those of allies and non-allied countries. The amount of trade using a lane, however, is weakly related to the additional sailing distance required to bypass a disruption. Across lanes with alternative routes, the correlation between these measures is very low. Some heavily used routes have nearby alternatives, while disruptions to less heavily used routes require substantial detours. These findings give empirical content to the distinction between exposure and vulnerability: trade value measures reliance on a lane, while maritime geography determines the alternatives available when access is restricted.

Countries also depend on foreign vessels to carry their trade, but where a ship is built need not determine who operates it. China and South Korea built close to four-fifths of active containership capacity, while carriers based in a broader set of countries operate the fleet. Roughly three-quarters of Chinese-built capacity is operated by non-Chinese carriers. A restriction on Chinese-built vessels therefore affects capacity operated by firms from many countries, not just Chinese carriers. We also find that a country's share of trade exposure on a lane is positively associated with capacity operated by its firms or provided by vessels built there, even on lanes between foreign ports.

Finally, we examine how trade exposure, limited rerouting options, and dependence on foreign vessels overlap across the shipping network. Some lanes that are relatively more important for Chinese than for U.S. trade are also difficult to bypass and heavily reliant on Chinese-built or Chinese-operated vessels. Dependence on these lanes therefore involves both the availability of alternative routes and access to vessel capacity. These findings show how dependence on maritime routes can be compounded by dependence on the vessels serving them.

To quantify the welfare consequences of maritime dependence, we study a spatial model of international trade extended to feature an endogenous shipping sector in which operators allocate vessel capacity across routes. Locations produce differentiated goods, households choose where to live within countries, and shipments can reach their destinations through multiple ports and alternative routes. Maritime transport requires vessel capacity, which operators allocate across shipping routes in response

to freight rates. A disruption therefore affects both the demand for shipping services and the capacity available to supply them. Freight rates adjust to clear shipping markets, while changes in transport costs affect trade, the location of production, and welfare.

The central feature of the shipping sector is that serving a lane requires a complete vessel rotation. Vessels must reach the loading port, carry the shipment, and return afterward. Longer voyages absorb more capacity, reducing the shipping services that a given fleet can provide. Operators can respond by reallocating vessels, but doing so changes the capacity available elsewhere in the network. A disruption therefore affects shipping costs beyond the routes directly exposed to it. This mechanism connects maritime geography with vessel availability: detours increase the capacity required to serve a route, while restrictions on vessels limit the capacity available to accommodate them.

We take the model to the data using the observed shipping network, bilateral trade, vessel deployments, and national income and population. Together, these observations determine the baseline allocation of economic activity, trade, and shipping capacity. We aggregate vessel capacity by operator country while retaining information on where each vessel was built, allowing us to distinguish restrictions on operators from restrictions based on construction origin. For each disruption that we consider, we solve for the equilibrium response of trade, routes, vessel deployment, and economic activity.

We evaluate the implications of the model in two ways. First, we investigate the extent to which the model's welfare costs line up with measured trade exposure. We increase the sailing requirement of each international lane by 10 percent, one at a time, so that differences in welfare costs reflect countries' exposure and the equilibrium adjustment rather than the severity of the disruption. Lanes that matter more for Chinese than for U.S. trade do impose larger welfare costs on China than on the United States. Yet measured exposure accounts for only 36.4 percent of the cross-lane variation in this difference. Much of the rest reflects the adjustment of producer prices and the spread of higher freight rates and port congestion to other lanes. Measured exposure thus captures the direct effect of a disruption, but its full welfare cost depends on how prices and the shipping network adjust.

Second, we use the Red Sea disruptions that followed the Houthi attacks on commercial vessels in late 2023 to validate the model's equilibrium response before turning to hypothetical disruptions. Holding structural parameters fixed, we introduce carrier-specific shocks proportional to observed withdrawals from Suez, scaled to match the aggregate decline in container capacity through the canal, and compare the remaining model predictions with observed vessel movements and U.S. import data. Our model reproduces the observed rerouting of shipping from Suez toward the Cape of Good Hope and accounts for most of the capacity that shifted between the two routes in the data. Across U.S. import markets, it also matches how cargo was rerouted away from Suez, and it predicts, as in the data, that freight charges rose more and cargo volumes fell more in markets that relied more heavily on Suez before the crisis. The model thus captures the direction and cross-market pattern of the observed responses.

Because the cost of a closure depends on the routes around it, disruptions that affect several neighboring chokepoints at once can be far more costly, as they remove the alternatives ships would otherwise

use. The waters around Taiwan and the South China Sea are a natural setting to study such disruptions, since several of the most heavily used chokepoints in the network, from the Taiwan Strait to the Straits of Malacca and Sunda, lie close to one another. There, closing a limited area around Taiwan lengthens voyages only modestly, but broader closures cut off ports and force much longer detours. Following scenarios studied in the policy literature (Lin et al., 2024*a*; Cancian, Cancian and Heginbotham, 2025), we next consider a selective disruption that restricts cargo around Taiwan while cargo to and from China continues to move. It imposes large losses on Taiwan and nearby economies, while China's welfare is nearly unchanged. We then examine how restrictions on China-related cargo at other straits in the region change the cost borne by China. Because this cargo can shift to nearby straits, China is worse off only when the Taiwan Strait, Malacca, and Sunda are restricted together, and these restrictions also affect third countries, some of which bear larger additional losses than China. Maritime geography thus allows disruptions to be selective, imposing large losses on some countries at little cost to others.

Finally, we study how trade depends on the vessels that sail these routes. Much of this capacity is built and operated abroad, so trade can be disrupted even when routes remain open, if operators withdraw their vessels or policy restricts their use. U.S. dependence on Chinese capacity, the focus of recent U.S. policy toward China's maritime and shipbuilding sectors, is a natural case to study. We first examine the consequences of such a withdrawal, making Chinese-operated or Chinese-built capacity largely unavailable on lanes serving the United States. The two restrictions affect different vessels, since most Chinese-built ships are operated by non-Chinese carriers, yet they impose similar costs on the United States. Because routes compete for the same fleet, the losses also spread abroad, and economies in Southeast and South Asia lose several times as much as the United States. We then ask whether policy can limit these losses by lowering reliance on this capacity in advance, focusing on the shipping fees announced by the United States in April 2025 and suspended in November 2025. Such fees work like insurance: they raise shipping costs in normal times but may lower the losses when a disruption occurs. The model implies that the fees lower U.S. welfare by 0.004 log percentage points in normal times and reduce the additional loss from a disruption, but by less than their cost during normal times.

This paper contributes to three literatures. First, it contributes to research on transport networks and spatial equilibrium, which studies how transportation costs and infrastructure shape trade, market access, and welfare (Allen and Arkolakis, 2014, 2022; Fajgelbaum and Schaal, 2020; Fan and Luo, 2020; Jaworski, Kitchens and Nigai, 2023; Bonadio, 2024; Fuchs and Wong, 2026). Within this literature, Fajgelbaum et al. (2026) study the design and international incidence of strategic Belt and Road infrastructure. Closer to our setting, Coşar and Thomas (2021) and Maurer, Milsom and Rauch (2026) study how maritime geography shapes routes and distances, and Feyrer (2021) uses the closure of the Suez Canal to identify the effect of sea distance on trade. Devetak, Verschuur and Klimek (2026) show in a vessel-level simulation that adaptive routing changes the losses from disruptions. Building on this work, we study how geographic restrictions interact with the supply of vessel capacity. Longer routes require additional capacity, and its reallocation changes shipping costs throughout the network. Incorporating

this adjustment allows us to distinguish countries' exposure to a route from the welfare consequences of disrupting it.

Second, we contribute to a growing literature studying shipping costs, vessel deployment, route structure, and round-trip imbalances (Kalouptsidi, 2014; Brancaccio, Kalouptsidi and Papageorgiou, 2020; Behrens and Picard, 2011; Heiland et al., 2025; Ganapati and Wong, 2023; Ganapati, Wong and Ziv, 2024; Wong, 2022; Dunn and Leibovici, 2024). Related work examines transshipment and liner networks (Coşar and Demir, 2018; Brooks, Gendron-Carrier and Rua, 2018; Notteboom and Rodrigue, 2008; Do et al., 2025), shipbuilding and vessel-specific policies (Kalouptsidi, 2018; Barwick, Kalouptsidi and Zahur, 2024, 2025; Chowdhry, Heiland and Mahlkow, 2026), and shipping disruptions and their consequences (Bai et al., 2026; Verschuur, Lumma and Hall, 2025; Reyes-Heroles and Zhou, 2026; Fuchs and Leibovici, 2026). Our contribution is to integrate endogenous shipment routing and vessel deployment into a tractable spatial trade model. Shipping supply accounts for the complete rotation required to serve a lane, and freight rates clear markets as trade and vessel capacity adjust across the network. We also distinguish operator country from construction origin, allowing restrictions on deployment control and vessel production origin to affect different parts of the fleet. A contemporaneous paper by Brockhaus, Hinz and Serfaty (2026) studies monopoly transit fees at maritime chokepoints. Our focus is instead on the welfare consequences of geographic disruptions, restrictions on vessel capacity, and policies intended to reduce maritime dependence.

Third, the paper contributes to research on geoeconomics and supply-chain resilience. Following Hirschman (1945), Keohane and Nye (1977), and Baldwin (1985), recent work studies how economic dependence creates political leverage (Farrell and Newman, 2019; Blackwill and Harris, 2016; Clayton, Maggiori and Schreger, 2024, 2026; Mohr and Trebesch, 2025). Related work measures indirect exposure to foreign suppliers through input–output linkages (Baldwin, Freeman and Theodorakopoulos, 2023) and shows how substitution shapes the losses from an energy cutoff (Bachmann et al., 2024). Research on resilience examines the costs and benefits of reducing foreign dependence through diversification, reshoring, and friend-shoring (Antràs, 2020; Antràs and Chor, 2022; Bonadio et al., 2021; Grossman, Helpman and Lhuillier, 2023; Javorcik et al., 2024; Adamopoulos and Leibovici, 2025), as well as the formation of supply networks and robust sourcing (Elliott, Golub and Leduc, 2022; Jiang, Rigobon and Rigobon, 2022). Whereas Kleinman et al. (2026) study bilateral capacity investment under uncertainty, we examine adjustment through existing routes and vessel capacity after a disruption. We bring maritime transportation into this analysis as a service on which international supply chains depend. Our framework distinguishes the welfare loss imposed on a country from the economic cost borne by a potential disruptor, while the policy analysis weighs the cost of reducing vessel dependence against the resulting change in disruption losses.

Section 2 develops the analytical framework linking maritime exposure, welfare, and geopolitical incentives. Section 3 presents the data and empirical findings, and Section 4 develops the quantitative model and describes its implementation. Section 5 examines standardized lane-level disruptions and

assesses the model against the Red Sea disruptions. Section 6 studies global and regional geographic vulnerabilities. Section 7 examines disruptions to vessel capacity and policies to reduce maritime dependence. Section 8 concludes.

2 Maritime Exposure and Vulnerability: Sufficient Statistics

This section develops an analytical framework to study how dependence on maritime trade maps into vulnerability to disruption. We first characterize how routed exposure, the availability of alternative routes and vessel capacity, and the adjustment of income and consumer prices determine the welfare consequences of a disruption. We then examine how these consequences influence the likelihood of geopolitical disruptions, taking into account the economic costs borne by both the affected country and the country that could initiate them.

2.1 Environment

We develop the analytical framework in a spatial Armington economy, following the spatial trade and transport-network literature (Allen and Arkolakis, 2014, 2022; Allen, Fuchs and Wong, 2025; Fuchs and Wong, 2026). The framework allows us to map changes in transport costs and producer prices to country welfare. In Section 4, we further specify the maritime routing and shipping supply and determine these changes in equilibrium.[2]

The world consists of M countries, indexed by $m \in \mathcal{C} = \{1, \ldots, M\}$. Each country contains a set of locations $\mathcal{N}_m$, and $\mathcal{N} = \bigcup_m \mathcal{N}_m$ denotes the global set of locations. Let o and d index origin and destination locations, with $m(o)$ denoting the country containing o. Location o has a mass L_o of workers who supply labor inelastically and earn wage w_o. Workers may move freely within their country but not across countries. Each country's total labor endowment is fixed.

Each location produces a differentiated good using labor, with technology $q_o = A_o L_o$, where A_o is local productivity. Production is competitive, implying a producer price $p_o = w_o/A_o$. Delivering one unit of this good to destination d requires τ_{od} units at origin, so its delivered price is $p_o \tau_{od}$. The delivery cost includes the shipping services required to transport the good through the maritime network.

Workers value the local amenity u_d and consume a CES aggregate of goods produced across locations, with elasticity of substitution $\sigma > 1$. They take prices and amenities as given and spend their labor income on consumption. Let P_d denote the consumption price index, $y_d = w_d L_d$ denote total expenditure at d, and X_{od} denote expenditure at d on the good produced at o. The corresponding expenditure share is $\lambda_{od} = X_{od}/y_d$.

A worker living at d obtains utility $u_d w_d / P_d$. In equilibrium, free mobility equalizes utility across

[2]Here, we hold local productivity and amenities fixed, so the analytical result applies to the benchmark without the local externalities allowed in that section.

populated locations within each country:

$$\frac{u_d w_d}{P_d} = W_m, \qquad d \in \mathcal{N}_m \text{ with } L_d > 0, \tag{2.1}$$

where W_m denotes country m's welfare level.

The goods market clearing condition equates the value of production at each location to expenditure on its good. Denoting R_o for production revenue at o, we obtain the following:

$$R_o \equiv \sum_d X_{od} = p_o q_o = w_o L_o = y_o, \tag{2.2}$$

where country m's income is $Y_m = \sum_{o \in \mathcal{N}_m} R_o$. Since households spend their labor income, each country's expenditure equals its income, so trade is balanced.

2.2 Tracing the Impact of Maritime Disruptions

Maritime disruptions affect welfare through transport costs and the prices at which countries buy and sell goods. We focus on three potential disruption scenarios here: geographic restrictions, shipping vessel withdrawals, and fees or access restrictions. Geographic restrictions can require longer voyages, vessel withdrawals can reduce available capacity for shipping, and fees or access restrictions can raise the cost of shipping services. Figure 2 illustrates these channels. We start by deriving the first-order welfare effects of disruptions as a function of baseline routed exposure and changes in equilibrium transport costs and producer prices. This identifies the economic determinants of vulnerability without requiring a particular specification of shipping supply.

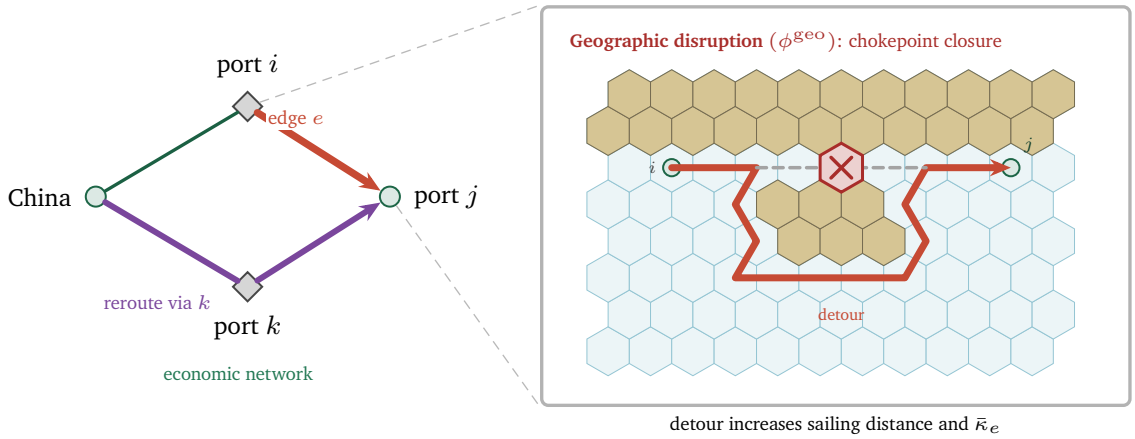
Let g denote a disruption scenario and d_g its intensity, with $d_g = 0$ at the baseline equilibrium. Throughout this subsection, derivatives and expenditure weights are evaluated at that equilibrium. The maritime network consists of directed shipping edges indexed by $e \in \mathcal{E}$. Let κ_e denote the generalized cost of using edge e, including freight rates and other geographic or port costs. We define its response to the disruption as

$$\phi_{e,g} \equiv \left. \frac{\partial \log \kappa_e}{\partial d_g} \right|_{d_g = 0}.$$

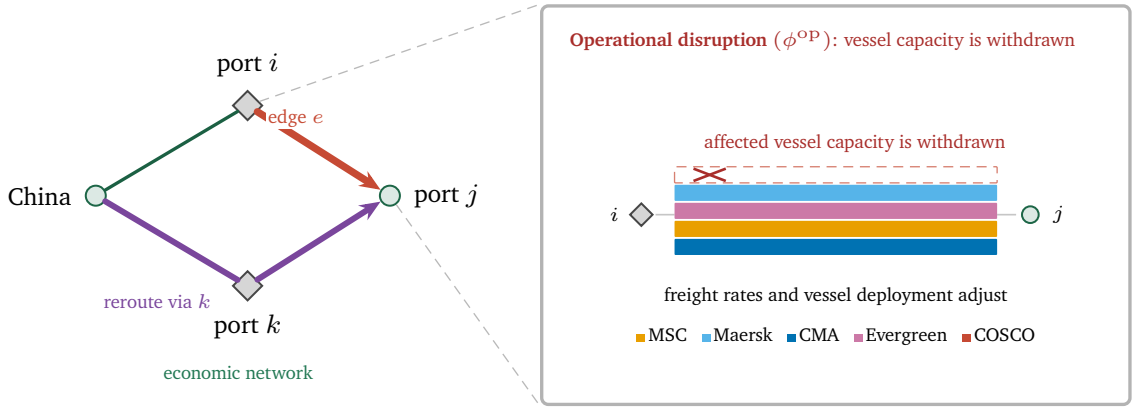
Next, we focus on the shipping carriers. Carriers choose which routes to use to deliver goods from origin o to destination d. Let x describe this routing choice, $\mathcal{R}_{od}$ denote the set of feasible choices, and $C_{od}(x; \kappa)$ be the cost of delivering one unit given edge costs $\kappa = \{\kappa_e\}_{e \in \mathcal{E}}$. Competitive carriers take edge costs as given and minimize delivery costs:

$$\tau_{od}(\kappa) = \min_{x \in \mathcal{R}_{od}} C_{od}(x; \kappa). \tag{2.3}$$

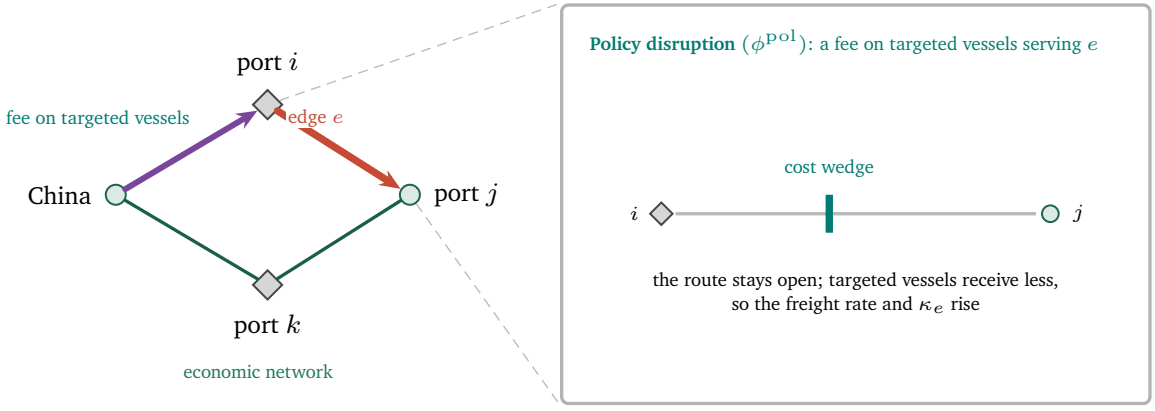
Thus, τ_{od} is the cost of the cheapest way to deliver goods between the two locations, allowing carriers to



(a) Geographic disruptions



(b) Disruptions to vessel capacity



(c) Policy interventions

Figure 2. Channels of maritime disruption. Each panel shows a disruption to the same shipping connection. Geographic restrictions require a detour, vessel withdrawals reduce available capacity, and policy interventions change the cost or conditions of access. Trade can also be redirected through an alternative connection.

change routes when edge costs change.

The effect of a cost increase on edge e depends on how much that edge contributes to the cost of delivering goods from o to d. We measure this contribution by

$$\upsilon_{od,e} \equiv \frac{\partial \log \tau_{od}}{\partial \log \kappa_e}. \tag{2.4}$$

A one percent increase in κ_e raises the delivery cost τ_{od} by approximately $\upsilon_{od,e}$ percent. Because carriers choose routes optimally, the envelope theorem allows us to calculate this first-order effect using their baseline routing choices. Changes in those choices do not add a separate first-order term. Depending on the transport technology, $\upsilon_{od,e}$ can be interpreted as a transport expenditure share or expected use of edge e.[3]

We can now combine the disruption's effect on edge costs from earlier with the routing weights to obtain its effect on delivery costs:

$$\frac{\partial \log \tau_{od}}{\partial d_g} = \sum_e \upsilon_{od,e} \phi_{e,g}. \tag{2.5}$$

Each term combines the change in the cost of using edge e, measured by $\phi_{e,g}$, with the sensitivity of delivery costs to that edge, measured by $\upsilon_{od,e}$. Summing across edges gives the total delivery-cost response. The same disruption can therefore affect shipments differently, depending on the routes they use.

A country's exposure to an edge depends on how much it trades and how heavily that trade relies on the edge. Let i and j index origin and destination countries. Total expenditure by country j on goods from country i is $X_{ij} = \sum_{o \in \mathcal{N}_i} \sum_{d \in \mathcal{N}_j} X_{od}$. We define this trade's exposure to edge e as

$$\Xi^X_{ij,e} = \sum_{o \in \mathcal{N}_i} \sum_{d \in \mathcal{N}_j} X_{od} \upsilon_{od,e}. \tag{2.6}$$

For each origin–destination pair, we multiply expenditure X_{od} by the sensitivity of its delivery cost to the cost of edge e. We then sum across locations in the two countries. Exposure is therefore greater when more trade depends on the edge or when its delivery cost is more sensitive to that edge's cost. If $\upsilon_{od,e}$ equals one when the shipment uses the edge and zero otherwise, $\Xi^X_{ij,e}$ is simply the value of trade from i to j passing through it. Summing over origins gives country j's routed exposure to edge e,

$$\Xi^X_{j,e} \equiv \sum_i \Xi^X_{ij,e}, \tag{2.7}$$

which Section 3 measures, along with export exposure, by assigning observed trade to routes.

Combining exposure with the change in edge costs gives the direct delivery-cost effect on country j's

[3] Appendix A.1 derives this result and both interpretations of $\upsilon_{od,e}$.

purchases from country i:

$$\Omega_{ij,g} = \sum_e \Xi^X_{ij,e} \phi_{e,g}. \tag{2.8}$$

We multiply baseline trade exposure on each edge by the disruption-induced change in that edge's cost, then sum across edges. Equivalently, $\Omega_{ij,g}$ weights each origin–destination pair's proportional delivery-cost response by its baseline expenditure. It therefore measures the direct cost of the disruption in value terms, before producer prices adjust. Summing $\Omega_{ij,g}$ across origins gives the corresponding effect on all goods purchased by country j, $\sum_i \Omega_{ij,g} = \sum_e \Xi^X_{j,e} \phi_{e,g}$.[4]

Our analysis in this section does not require specifying the supply of shipping services: the results below hold for any process that determines $\phi_{e,g}$. Different disruptions, however, affect supply through different margins. Closures lengthen voyages, capacity withdrawals remove vessels from some uses, and fees drive a wedge between the price shippers pay and the price operators receive. Section 4 models these margins explicitly: operators allocate vessels across complete rotations, and freight rates clear the market for shipping services on each edge, determining $\phi_{e,g}$ in equilibrium.

We now connect these delivery-cost changes to household welfare. Households pay the delivered price $p_o \tau_{od}$ for goods from origin o, so their cost of living depends on both transport costs and producer prices. Under CES preferences,

$$\mathrm{d} \log P_d = \sum_o \lambda_{od} \left(\mathrm{d} \log p_o + \mathrm{d} \log \tau_{od} \right). \tag{2.9}$$

Each delivered-price change affects the price index in proportion to the good's baseline expenditure share. These shares suffice for the first-order calculation because households choose consumption optimally.

Producer prices are also equilibrium outcomes and affect both the expenditure and income sides of welfare. They adjust to clear goods markets as the disruption changes demand across origins. Higher prices for the goods a country purchases raise its cost of living, while higher prices for the goods it produces raise labor income. The latter follows from competitive production and fixed productivity, which imply that wages and producer prices change in the same proportion. Let $r_{o,g} = \partial \log p_o / \partial d_g$ denote the producer-price response at location o.

Let $\ell_{m,g} = -\partial \log W_m / \partial d_g$ denote the marginal proportional welfare loss. Define the marginal equivalent-variation (EV) loss by $\mathcal{W}_{m,g} = Y_m \ell_{m,g}$, with positive values indicating losses. EV values the welfare change at baseline prices. With fixed amenities and homothetic preferences, aggregating EV at their baseline locations gives the finite loss $Y^0_m(1 - \widehat{W}_m)$; its derivative at the baseline is $\mathcal{W}_{m,g}$. The baseline-income-weighted log loss agrees with EV only to first order. Combining the effects on consumer prices and labor income gives the following result.

[4] Selective disruptions, such as restrictions on particular cargo classes, change generalized edge costs with traffic-specific equilibrium elasticities $\phi_{od,e,g}$. Exposure then enters in pair form, $\Omega_{ij,g} = \sum_e \Xi^X_{ij,e} \phi_{ij,e,g}$, where $\phi_{ij,e,g}$ is the $\Xi^X_{ij,e}$-weighted average of shipment-level elasticities on edge e for the country pair (Appendix A.1). The disruption around Taiwan and the restrictions on China-related cargo in Section 6.2 are selective in this sense.

Proposition 2.1 (Welfare effects of a maritime disruption)**.** *Consider the Armington benchmark described above, with fixed local productivity and amenities, labor-income spending, and goods-market clearing. Assume that equilibrium outcomes, including edge costs and optimized routing costs, are differentiable at the baseline along disruption scenario g.*

(i) Country j's marginal equivalent-variation loss is

$$\mathcal{W}_{j,g} = \underbrace{\sum_{e} \Xi^{X}_{j,e}\, \phi_{e,g}}_{\text{routed exposure} \times \text{disruption response}} + \underbrace{\mathcal{T}^{\text{ToT}}_{j,g}}_{\text{producer-price incidence}}, \tag{2.10}$$

where

$$\mathcal{T}^{\text{ToT}}_{j,g} \equiv \sum_{d \in \mathcal{N}_j} \sum_{o} X_{od} r_{o,g} - \sum_{o \in \mathcal{N}_j} R_o r_{o,g}.$$

(ii) If producer prices change in the same proportion across locations within each country, let $r_{i,g}$ denote the common response in country i. Producer-price incidence for country j becomes

$$\mathcal{T}^{\text{ToT}}_{j,g} = \sum_{i} X_{ij} r_{i,g} - Y_j r_{j,g}.$$

The proposition separates the welfare effect into delivery costs and producer-price adjustment. The first term values the change in delivery costs using baseline expenditures. The second accounts for changes in the prices of the goods a country purchases and produces. Higher producer prices raise consumption costs for buyers but labor income for producers. Their net effect can therefore amplify or offset the loss from higher delivery costs.

Routed exposure is consequently insufficient to determine welfare losses. Exposure measures how much trade depends on an edge. The availability of alternatives and the response of shipping supply determine how its cost changes following a disruption. Producer-price adjustment then affects the distribution of losses across countries. Countries with similar exposure can therefore fare very differently, and some may even gain, as demand shifts toward their goods and raises the prices they receive.

The result characterizes the first-order welfare effect at the baseline equilibrium. Optimal consumption and routing choices allow this effect to be evaluated using baseline weights, as in the production-network and trade-welfare literature (Hulten, 1978; Baqaee and Farhi, 2019; Arkolakis, Costinot and Rodríguez-Clare, 2012; Costinot and Rodríguez-Clare, 2014). Larger disruptions require accounting for changes in those weights. Section 4 develops the model used to compute these counterfactuals, allowing trade, routes, vessel deployment, and prices to adjust jointly. Appendix A.2 provides the derivation and discusses the assumptions underlying the welfare result.

The quantitative specification also permits productivity and amenities to respond to local population. Appendix A.5 states how these features modify the fixed-productivity, fixed-amenity benchmark above, and distinguishes them from finite changes in baseline weights.

2.3 Geopolitical Incentives and Maritime Disruptions

The section before characterized the welfare consequences of a given disruption, and we now examine how those consequences shape the incentives to initiate it. A country may value imposing losses on another country while bearing economic costs itself, so it carries out the disruption only when the loss it can inflict outweighs its own cost, and otherwise it would leave maritime access unchanged.

Let a denote the country considering a disruption and h the target country. Country a chooses between leaving maritime access open, denoted by O, and carrying out disruption g. For each of the two scenarios, let $W_m(\cdot)$ denote country m's welfare level in the resulting equilibrium, as defined in equation (2.1). As such, $W_m(O)$ is the baseline level and $W_m(g)$ is the level after the full disruption, and their difference is the finite counterpart of the marginal loss in Proposition 2.1.

Country a cares about two things. The first is its own material payoff $V_a(\cdot)$, which comprises its real income and net shipping earnings under the two scenarios, less any direct costs of implementing the disruption. The second is the loss imposed on target country h, which country a values at weight χ. This weight is drawn from distribution F_χ and observed by country a before it chooses. Country a chooses the alternative that yields the higher utility U_a, where

$$U_a(O;\chi) = V_a(O), \qquad U_a(g;\chi) = V_a(g) + \chi Y_h \left[\log W_h(O) - \log W_h(g)\right]. \tag{2.11}$$

The second term in $U_a(g;\chi)$ measures the value country a assigns to reducing country h's welfare. Target country h's log welfare loss is multiplied by its baseline income Y_h, and χ determines how much that loss matters relative to the disruptor's own material payoff $V_a(g)$. A larger χ therefore implies a greater willingness to bear economic costs to impose losses abroad. This formulation follows the economic-statecraft and sanctions literature (Schelling, 1966; Eaton and Engers, 1992; Becko, 2024; Bianchi and Sosa-Padilla, 2024). We take this willingness as given rather than model the conflict behind it.

The decision depends on the loss imposed on the target and the cost borne by the disruptor. We measure both relative to open access:

$$L^{\text{fin}}_{h,g} = Y_h \left[\log W_h(O) - \log W_h(g)\right], \qquad C^{\text{fin}}_{a,g} = V_a(O) - V_a(g).$$

Here, $L^{\text{fin}}_{h,g}$ is the target's log welfare loss weighted by baseline income, and $C^{\text{fin}}_{a,g}$ is the reduction in the disruptor's material payoff. If the disruptor bears only its own welfare loss, its cost follows from Proposition 2.1. Otherwise, it also includes any shipping-income losses or implementation costs not already counted in welfare.

Country a prefers disruption when $U_a(g;\chi) \geq U_a(O;\chi)$ (Equation 2.11), which is equivalent to $\chi L^{\text{fin}}_{h,g} \geq C^{\text{fin}}_{a,g}$, so the choice reduces to a threshold on χ stated in the following proposition (See Appendix A.4 for proof):

Proposition 2.2 (Disruption incentives)**.** *Consider country a's decision to initiate disruption g against country h. If $L^{\text{fin}}_{h,g} > 0$ and $C^{\text{fin}}_{a,g} > 0$, country a weakly prefers disruption to open access if and only if its willingness χ reaches the geopolitical threshold*

$$\chi \geq \chi^{\text{fin}}_{a,h,g} \equiv \frac{C^{\text{fin}}_{a,g}}{L^{\text{fin}}_{h,g}}. \tag{2.12}$$

A country's vulnerability depends not only on the losses a disruption would impose on it, but also on how costly that disruption would be for the country initiating it. A large potential loss may offer little incentive to disrupt shipping if the disruptor would also bear substantial costs.

For the choice between disruption g and open access, the threshold also determines disruption risk. When target losses and disruptor costs are positive and F_χ is continuous, the probability that country a disrupts is $1 - F_\chi(\chi^{\text{fin}}_{a,h,g})$. Holding this distribution fixed, a larger target loss relative to the disruptor's cost lowers the threshold and weakly increases the probability of disruption.

We show that maritime dependence can create opportunities to impose losses abroad (Hirschman, 1945; Farrell and Newman, 2019), while the disruptor's own dependence can discourage such actions (Barbieri, 1996; Martin, Mayer and Thoenig, 2008). The threshold in equation (2.12) compares each disruption with open access. To choose among disruptions, we compare their net payoffs, $\chi L^{\text{fin}}_{h,g} - C^{\text{fin}}_{a,g}$, rather than their benefit-cost ratios alone. In the remainder of the paper, we measure the trade and shipping relationships underlying these losses and quantify their equilibrium consequences. Section 3 presents the data, and Section 4 develops the model used in the counterfactual exercises.

3 Data and Measurement

This section measures the sources of maritime vulnerability identified by the framework in the previous section. We document how exposed countries' trade is to shipping routes, how much room they have to reroute around disruptions, and how heavily they rely on foreign-operated and foreign-built vessels. Together, these measures distinguish reliance on maritime transportation from the geographic and supply conditions that make disruptions costly. They also provide the empirical foundation for our subsequent equilibrium analysis of how disruptions affect shipping costs and welfare across countries.

3.1 Data: Shipping and Trade Network

This section describes the data we use to trace countries' trade through the container shipping network. We combine three sources: vessel registries and scheduled port calls, bilateral trade data, and bill of lading data.

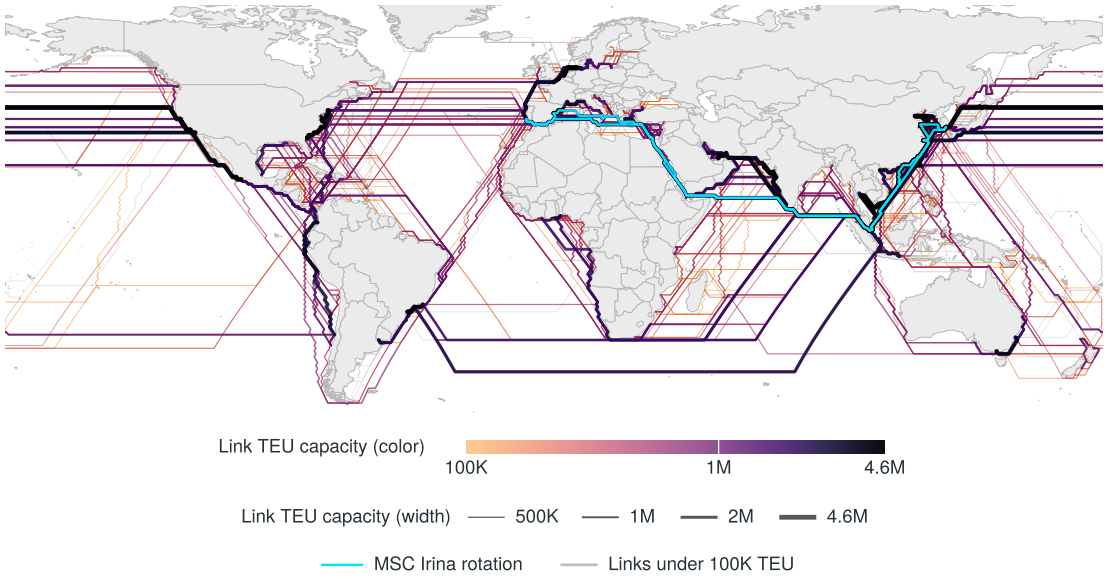


Figure 3. Directed containership network. Lanes connect successive port calls of scheduled liner services and are drawn along their sea paths. Color and width indicate annual container capacity, on a log scale for color and a square-root scale for width, both from 100,000 TEU to the 99th percentile (truncated above). Lanes with less than 100,000 TEU are shown in gray. The highlighted route is the rotation of the *MSC Irina* (MSC's Jade service). Source: vessel registry and rotation data.

Vessels and rotations data. Container shipping often takes place indirectly through intermediate hub ports, so a country's trade can depend on connections far from its own shores. To identify these connections and measure the vessel capacity serving them, we combine vessel registries with records of scheduled port calls. For each vessel, we observe its unique identifier, container capacity, construction year, builder firm and country, operating firm and country, and sequence of port calls. Capacity is measured in twenty-foot equivalent units (TEU). Successive port calls define directed shipping lanes. Vessel identifiers connect these lanes to the characteristics of the ships serving them.

To measure annual capacity, we account for both vessel size and service frequency. Annual rotation frequency is not directly observed. We impute it from the time required to complete a rotation, using sailing distance, vessel speed, and port operational assumptions. Let k_v denote vessel v's TEU capacity, n_v its imputed annual rotation frequency, and $m_{v,e}$ the number of times it traverses lane e during one rotation. Total annual container capacity on lane e, in TEU per year, is then the sum of $k_v n_v m_{v,e}$ across all vessels that serve the lane:

$$T_e = \sum_v k_v n_v m_{v,e}. \tag{3.1}$$

The schedules record 1,606 liner services, run by 5,245 vessels and calling at 768 ports in 166 countries and territories. Figure 3 maps the network and highlights the rotation of the *MSC Irina*, one of the world's largest containerships, with approximately 24,000 TEU of capacity. The example illustrates how a single vessel supplies capacity to several port-to-port connections within a complete rotation.

Trade data. We use 2023 product-level trade values and quantities from BACI (Gaulier and Zignago, 2010). To convert them into containerized sea trade, we apply predicted sea shares by country pair and industry (Verschuur, Koks and Hall, 2022), adjusted for products that move mostly by air, together with U.S. Census shares of containerized vessel trade (Census Bureau, 2023–2024).

Bill of lading data. The bill of lading records from Panjiva (S&P Global Market Intelligence, 2023–2024) are used in three ways. First, they provide information on the container payloads, which we use to convert trade quantities into TEU. Second, for U.S. imports they record the foreign port of lading and the U.S. port of entry of each shipment, which we use to validate the predictions from our routing network described below. Third, we match this data to our vessel rotation data to observe how U.S.-bound cargo was rerouted during the Red Sea crisis, which provides a separate validation of our model's predictions for counterfactual changes (Section 5.2).

Assigning trade through the container shipping network. While the bilateral trade data identify where goods are exported and imported between origin and destination countries, and the shipping rotation data identify the connections available between them, neither source alone reveals the journey taken by a shipment. In order to trace countries' trade through the container shipping network, we take the total 161.7 million TEU of demand and allocate it recursively through the shipping network (Allen, Fuchs and Wong, 2025; Fuchs and Wong, 2026). Conditional on each origin-destination pair, the allocation distributes every flow across direct and indirect journeys on the scheduled services in the vessel rotations data. One related approach is Heiland et al. (2025), who use AIS ship tracking data to compute the fastest route between port pairs in their empirical analysis and let shipments use multiple routes in their quantitative model. Our assignment likewise spreads each flow over several routes and, in addition, respects the capacity of each lane.

For each lane, we can calculate the value of each country-pair's trade that passes through it. Let X_{odk} be the value of product k shipped from region o to region d, and let $\upsilon_{odk,e}$ be the expected number of times this flow goes through lane e. The value of trade from country a to country b on lane e is

$$T_{e,ab} = \sum_{o \in a,\, d \in b,\, k} X_{odk}\, \upsilon_{odk,e}. \tag{3.2}$$

A transshipped shipment keeps its origin and destination but is counted on every lane it passes through, so summing across lanes measures shipping activity (in value) rather than its end-to-end trade.

Capacity, transshipment, and validation. We require the trade assignment process to meet two conditions. First, each lane carries at most its annual TEU capacity T_e from Equation (3.1). When a lane fills up, cargo moves to other routes, while bilateral trade and the vessel schedules stay fixed. In our data, the observed capacity is enough to transport all observed demand. Second, we set the cost of

changing ships at a hub so that transshipment makes up the share of world port handling reported by UNCTAD (2023), about 27 percent.

Next, we validate this assignment against independent data. For U.S. imports, we find high positive correlations between the trade assignment and the Panjiva bills of lading for July to November 2023, at 0.86 across foreign ports of lading and 0.96 across East and Gulf Coast ports of entry. We use these lane-level trade values to measure trade exposure in the next subsection, and they are also used as the baseline of our quantitative model.

3.2 Absolute and Relative Trade Exposure

We now use the assigned trade flows from the previous subsection to measure each country's trade exposure to individual lanes, how it differs across countries, and how much of it lies in foreign waters.

Measurement. Let $T_{e,ab}$ denote the value of trade from country a to country b that uses lane e (Equation (3.2)). Country m's trade exposure is the total value of its exports and imports using that lane:

$$\widetilde{\Xi}^{X}_{m,e} = \underbrace{\sum_{b} T_{e,mb}}_{\text{export exposure}} + \underbrace{\sum_{a} T_{e,am}}_{\text{import exposure}} . \tag{3.3}$$

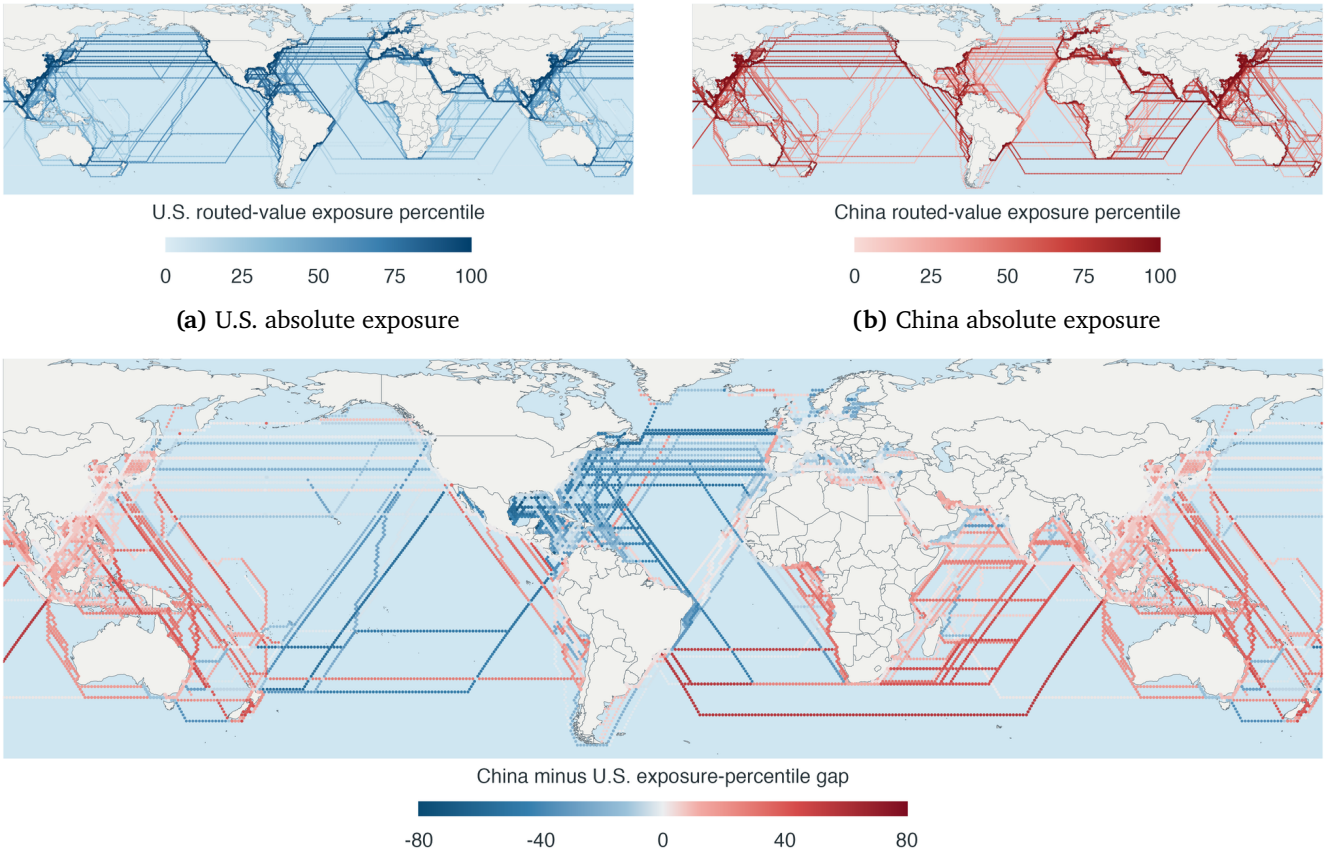
A shipment contributes to every lane along its route, so exposure across lanes can add up to more than total trade. As in Section 2, we weight lanes by the value of the trade they carry. While we combine imports and exports here to measure overall dependence on the network, the welfare calculation in later sections treats their effects on expenditure and income separately.

To compare the relative importance of a lane across countries, we rank lanes by exposure within each country. For countries H and F, the difference in percentile ranks is

$$\Delta\Xi^{\text{pct}}_{H-F,e} = \text{Pct}_H\left(\widetilde{\Xi}^{X}_{H,e}\right) - \text{Pct}_F\left(\widetilde{\Xi}^{X}_{F,e}\right). \tag{3.4}$$

A positive value means that the lane ranks higher in country H's trade network than in country F's. We refer to this relative dependence as the relative trade exposure. In share form, the relative trade exposure is the difference between two countries' shares of routed value carried by the lane, $\widetilde{\Xi}^{X}_{m,e} / \sum_{e'} \widetilde{\Xi}^{X}_{m,e'}$. For China and the United States it is $\Delta\widetilde{s}^{X}_{e}$, China's share minus the U.S. share in percentage points. This measure is an input to the comparison of target losses and disruptor costs in Section 2; as that section shows, whether disrupting a lane provides economic leverage also depends on the availability of alternatives and the costs imposed on both countries.

Exposure across countries. Panels (a) and (b) of Figure 4 report the within-country exposure ranks for the United States and China. U.S. exposure is concentrated on trans-Pacific and North Atlantic routes



(a) U.S. absolute exposure

(b) China absolute exposure

(c) China minus U.S. exposure-percentile gap

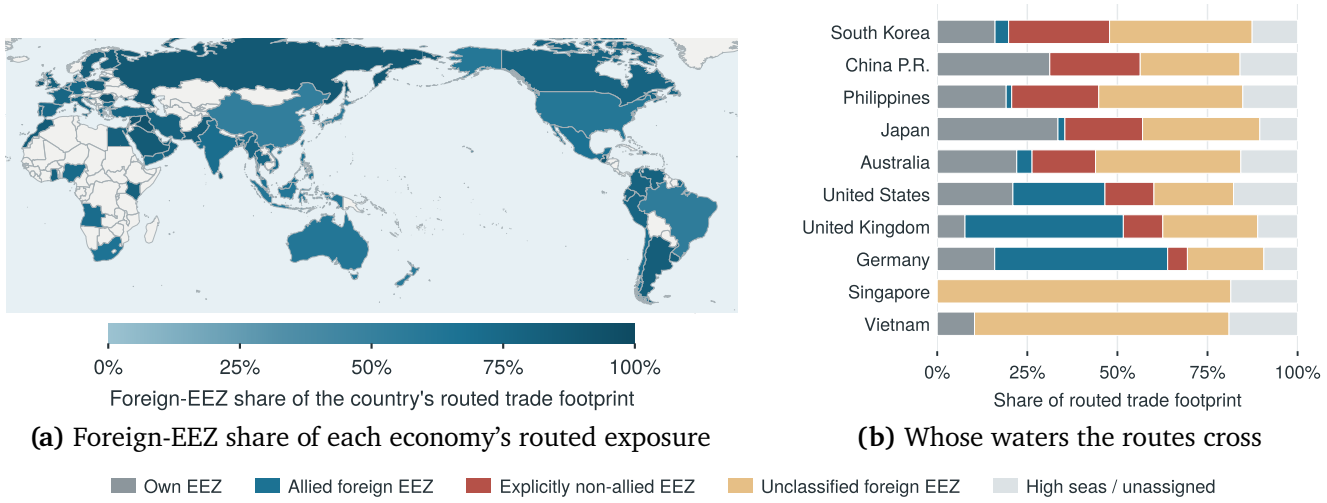
Figure 4. Absolute and relative trade exposure. Panels (a) and (b) rank lanes by their assigned import and export value within each country. Panel (c) reports the China-minus-U.S. difference in these percentile ranks: blue indicates greater relative importance for U.S. trade and red for Chinese trade. Source: bilateral trade and the directed shipping network.

and the connections serving those gateways. Chinese exposure is concentrated on intra-Asian hubs and Asia–Europe routes. Both countries depend on connections outside their own port networks because their imports and exports travel through foreign hubs.

Panel (c) reports the China-minus-U.S. difference in exposure ranks. It identifies lanes that are relatively more important to one country's trade than to the other's. A lane can carry substantial trade for both countries yet rank much higher for one than for the other.

Stylized Fact 1 (Exposure extends beyond domestic ports). Countries depend on foreign gateways and connections between third countries, as well as the lanes serving their own ports.

Dependence on foreign maritime geography. Next, we ask how much of a country's trade-exposed shipping lanes travels through foreign waters. To calculate this, we map the shipping lanes onto exclusive economic zone (EEZ) boundaries from Flanders Marine Institute (2023). EEZs are the waters within 200 nautical miles of a country's coast. For each lane, we calculate the fraction of its route inside foreign EEZs and average these fractions using the country's trade exposure on each lane as weights. This measure



(a) Foreign-EEZ share of each economy's routed exposure

(b) Whose waters the routes cross

Figure 5. Dependence on foreign maritime geography. Panel (a) reports the exposure-weighted foreign-EEZ share. Panel (b) separates the exposure of selected countries into own, allied, explicitly non-allied, unclassified foreign, and high-seas or unassigned exposure; the categories sum to one. Source: assigned bilateral trade and Marine Regions EEZ v12.

captures the foreign-water share of the routes serving each country's trade. Every country crosses some foreign waters to trade, so the measure differs through the weights. An Asia-Europe lane through the Red Sea and the Mediterranean passes along the coasts of many countries, while a trans-Pacific crossing spends much of its length on the high seas. A country with more exposure on the first kind of lane therefore has a higher foreign-water share.

Figure 5a maps the foreign-EEZ share of each country's trade exposure. The median share is 78 percent. The share is 61 percent for the United States and 53 percent for China. Weighting countries by their total lane-level trade exposure, we calculate an aggregate share of 68 percent.

Figure 5b then distinguishes whose waters these routes cross: the country's own EEZ, those of its allies, those of non-allied countries, other foreign EEZs, and waters outside the EEZ classification. We define non-allied narrowly: China and the members of the formal U.S. alliance coalition count as non-allied with each other, and all other countries remain unclassified. Foreign-water exposure depends on where a country is geographically. European countries such as the United Kingdom and Germany rely heavily on allied waters, while a substantial share of Japan's and South Korea's routes crosses China's EEZ. The United States falls in between, with a sizable allied share and a large share of high seas on trans-Pacific routes. Its implications for vulnerability therefore depend on which country could impose a restriction and the economic costs it would face.

Pair-level dependence and trade. We can compute the same measure for each exporter–importer pair. For each pair, we split its routed container value by service leg and record whose EEZs each leg crosses. Summing these pair measures over a country's partners gives back the country shares in Figure 5. We find that third countries' waters make up a large part of most routes: for the median pair, 80 percent of

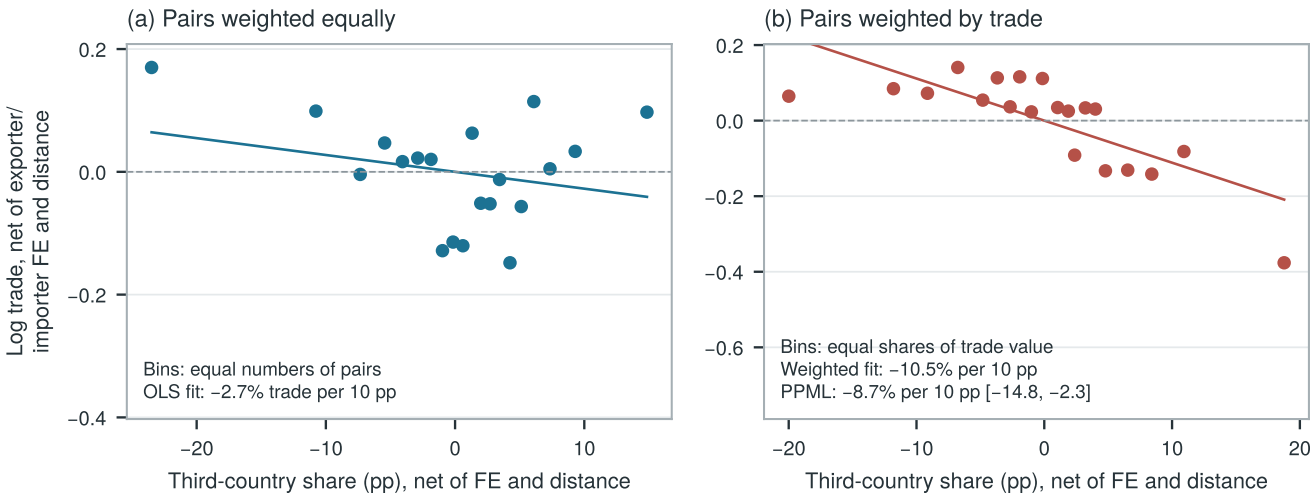


Figure 6. Third-country dependence and bilateral trade. Binned scatter of log bilateral container trade against the pair's share of route cells in third countries' waters, both net of exporter and importer fixed effects and log distance; 20 bins. Panel (a) weights pairs equally, panel (b) by trade value; lines are the corresponding linear fits, and panel (b) also reports the PPML gravity coefficient. Source: BACI, the capacity-constrained assignment, and Marine Regions EEZ v12.

the route lies in the waters of third countries, and 62 percent lies in the waters of countries allied with neither partner. Weighted by trade value, the third-country share is 48 percent.

Figure 6 relates this dependence to bilateral trade. Pairs whose trade routes run more through third countries' waters trade less, especially among large trade flows. We estimate a PPML gravity regression with exporter and importer fixed effects, controlling for distance, contiguity, common language, and colonial ties. A 10 percentage point higher third-country share is correlated with 8.7 percent less trade (95 percent interval: 2.3 to 14.8 percent). The estimate is 8.2 percent when we also control for the partners' distance in UN General Assembly voting. When we weight pairs equally, the relation is much weaker (panel a).

Stylized Fact 2 (Foreign dependence and non-allied exposure are distinct). Countries' trade relies extensively on routes through foreign waters, including those of allies and non-allied countries. Vulnerability depends on which country can restrict access, the trade exposed to that restriction, and the costs the country would bear by imposing it.

3.3 Geographic Vulnerability

The geography of the shipping network determines what alternatives are available when a route is disrupted, so the amount of trade on a route need not reflect the cost of losing access to it. In this subsection we measure the extra sailing distance needed to bypass local disruptions across the network and examine how it relates to trade exposure.

Measurement. We represent navigable waters as a network of hexagonal cells approximately 100 kilometers across and identify the least-cost sailing route between each pair of connected ports. We then consider a disruption to one cell at a time, requiring ships to avoid that cell, and calculate an alternative route between the same ports. We exclude cells containing the origin or destination port so that the exercise measures rerouting rather than port closures. If the ports cannot be connected through the remaining network, we record that outcome separately from cases in which ships can take a longer route.

Let d_r^0 denote the sailing distance between the ports on route r before the disruption and d_r^g the distance required to connect them while avoiding the disrupted cell. When an alternative route exists, the proportional increase in sailing distance is

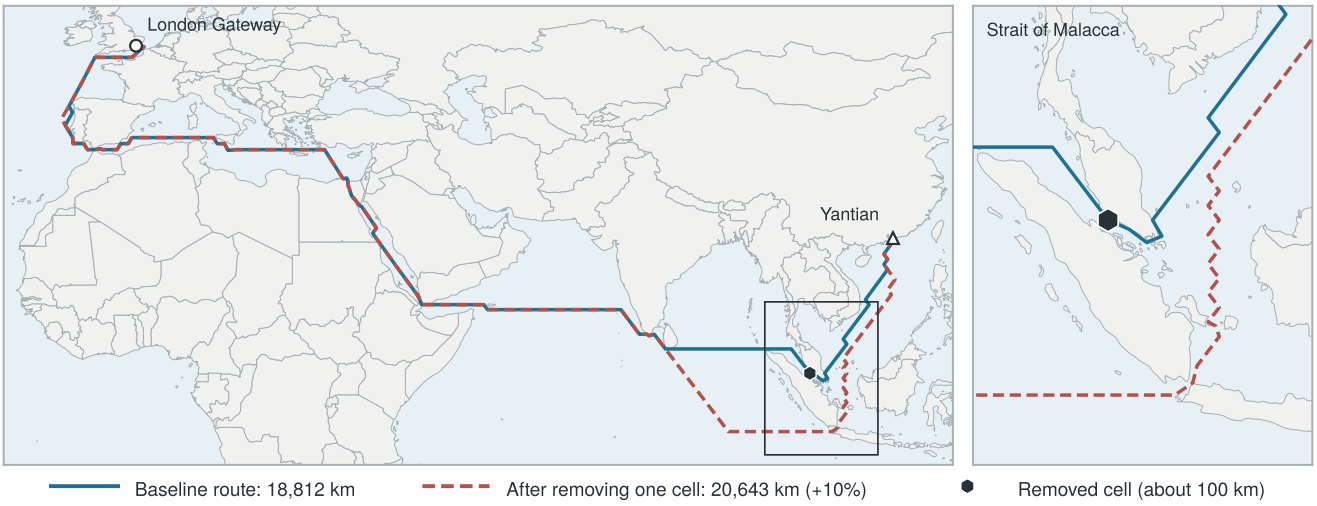
$$\Delta_{r,g}^{\text{geo}} = \frac{d_r^g - d_r^0}{d_r^0}. \tag{3.5}$$

Each experiment measures the detour around one disrupted cell, rather than the consequences of closing an entire strait. Where a passage is narrower than the grid, however, a single cell can close it.

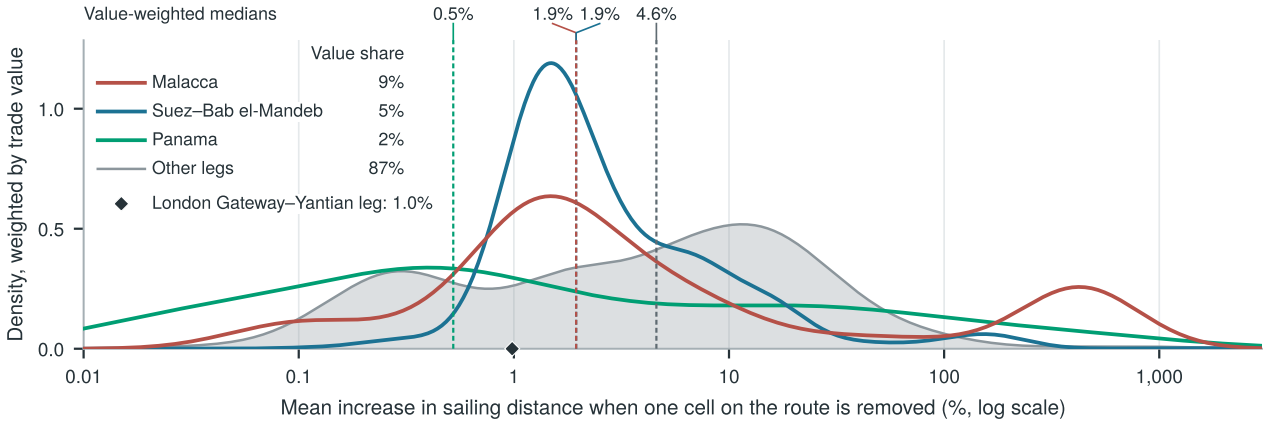
As in Coşar and Thomas (2021), changes in sailing distance measure the geographic consequences of a disruption. Longer voyages raise transport costs and require more vessel capacity to serve the same ports. The resulting change in shipping costs also depends on how trade and vessel deployment adjust. The distance calculations therefore measure a determinant of the cost response $\phi_{e,g}$ in Section 2, rather than that response itself. Section 4 determines these adjustments in equilibrium.

Trade exposure and detours. Panel (a) of Figure 7 illustrates how a local disruption near Malacca changes the sailing route between two ports. Avoiding one ocean cell in the strait sends the London Gateway–Yantian service south of Sumatra and through the Sunda Strait, which lengthens the voyage from 18,812 to 20,643 km, or 10 percent. Most other cells on the same route can be avoided at little cost: averaged over all of its cells, the increase is 1.0 percent. Panel (b) compares the estimated trade value carried by each lane with the average proportional increase in sailing distance required to bypass local disruptions. The comparison includes every lane that some local disruption lengthens, which carry 78 percent of routed value. For the value-weighted median lane, a disruption adds 1.9 percent to sailing distance on lanes through Malacca, 1.9 percent through Suez–Bab el-Mandeb, 4.6 percent elsewhere, and 0.5 percent through Panama. The correlation between the two measures in logs is -0.032 across these lanes outside the extreme tails of either measure. Lanes carrying more trade therefore do not generally require larger detours when disrupted. This finding illustrates the distinction in Section 2: exposure measures reliance on a route, while the availability of alternatives helps determine the cost of losing access to it. Section 6 evaluates the resulting welfare effects, accounting for longer voyages, vessel reallocation, and changes in trade and prices.

Stylized Fact 3 (Geographic Exposure and the cost of rerouting). Trade value and the additional sailing



(a) Local Malacca disruption and rerouting



(b) Routed trade value and bypass difficulty

Figure 7. Exposure and route substitution. Panel (a) shows the baseline path and detour of the London Gateway–Yantian leg after a local disruption near Malacca. Panel (b) compares trade value with the mean increase in sailing distance, highlighting chokepoint groups. The mean is taken over each leg's bypassable cells (value-weighted densities, log scale; dashed value-weighted medians 1.9, 1.9, 0.5 and 4.6 percent for Malacca, Suez–Bab el-Mandeb, Panama and other legs). The three chokepoint groups overlap, since a leg through two chokepoints belongs to both, and other legs pass none of them; routed value counts a shipment once on every leg it traverses. The experiments remove individual ocean cells, not entire straits. Source: bilateral trade assignment and the ocean network.

distance required to bypass a disruption are weakly related. The routes carrying the most trade are not necessarily those that are most difficult to replace.

3.4 Vulnerability to Foreign Fleets

Access to maritime routes is only one source of shipping dependence. Countries also rely on the vessels that carry their trade. We study this dependence on foreign fleets by separating where vessels were built from where their operators are based. We look at both because the two can differ and matter separately. A Chinese-built ship, for example, may be run by a firm from another country. Restrictions on Chinese operators and restrictions on Chinese-built ships therefore affect different parts of the fleet, and they can have different consequences for shipping and trade. This distinction is also relevant to the shipping fees announced by the United States in April 2025, which distinguish vessels by operator and by builder origin (USTR, 2025).

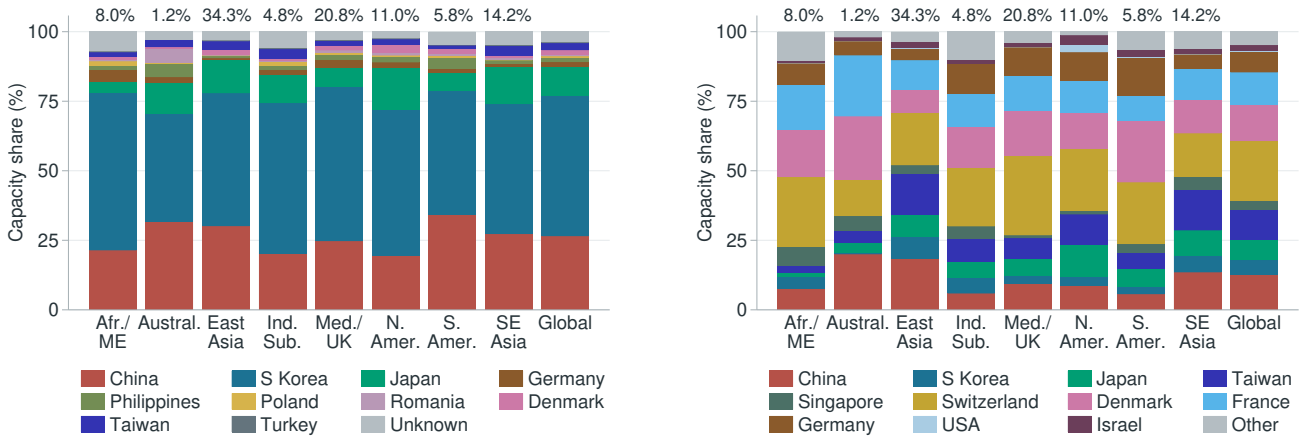
For each lane, we measure the share of annual shipping capacity operated by firms from country F, $s^{\text{op}}_{F,e}$, and the share provided by vessels built in F, $s^{\text{build}}_{F,e}$. These shares account for vessel size and service frequency, using the same weights as annual lane capacity T_e (Equation (3.1)). We report the composition of the global fleet in Figure 8, weighting each vessel by its annual capacity at ports in each region.

Concentration and vessel origin. Figure 8 compares the geography of shipbuilding and vessel operation. Panel (a) shows that construction is concentrated in East Asia: China and South Korea together account for close to four-fifths of active containership capacity. Panel (b) shows a more dispersed distribution across operator countries. At the carrier level, the largest firm accounts for roughly one-fifth of capacity and the three largest for about one-half. Roughly three-quarters of Chinese-built capacity is operated by non-Chinese carriers.

The fleet's composition reflects investment over many years: containerships commonly remain in service for 20–30 years (Stopford, 2008). China and South Korea have built close to nine-tenths of new capacity since 2020, while East Asian carriers account for a larger share of newer vessels than of older cohorts. China's rising share reflects the shipbuilding expansion documented by Barwick, Kalouptsidi and Zahur (2024, 2025).

Reducing dependence on this capacity can also entail costs. Larger vessels spread capital and operating costs over more containers, and orderbook prices per TEU decline with vessel size (Stopford, 2008). These patterns motivate the policy trade-off studied in Section 7, although they do not measure the cost of replacing a country's shipbuilding capacity.

Capacity and country trade. We examine whether a country's imports and exports account for a larger share of trade exposure on lanes with more capacity operated by its firms or supplied by vessels built there. Controlling for lane and country fixed effects, a 10 percentage-point increase in the country's



(a) Builder-country shares

(b) Operator-country shares

Figure 8. Builder and operator concentration. Panels report country shares of annual containership capacity at ports in each region, classified by builder and operator country. Annual capacity is vessel TEU times trips per year; the number above each bar is the region's share of global capacity. Source: vessel registries and liner schedules.

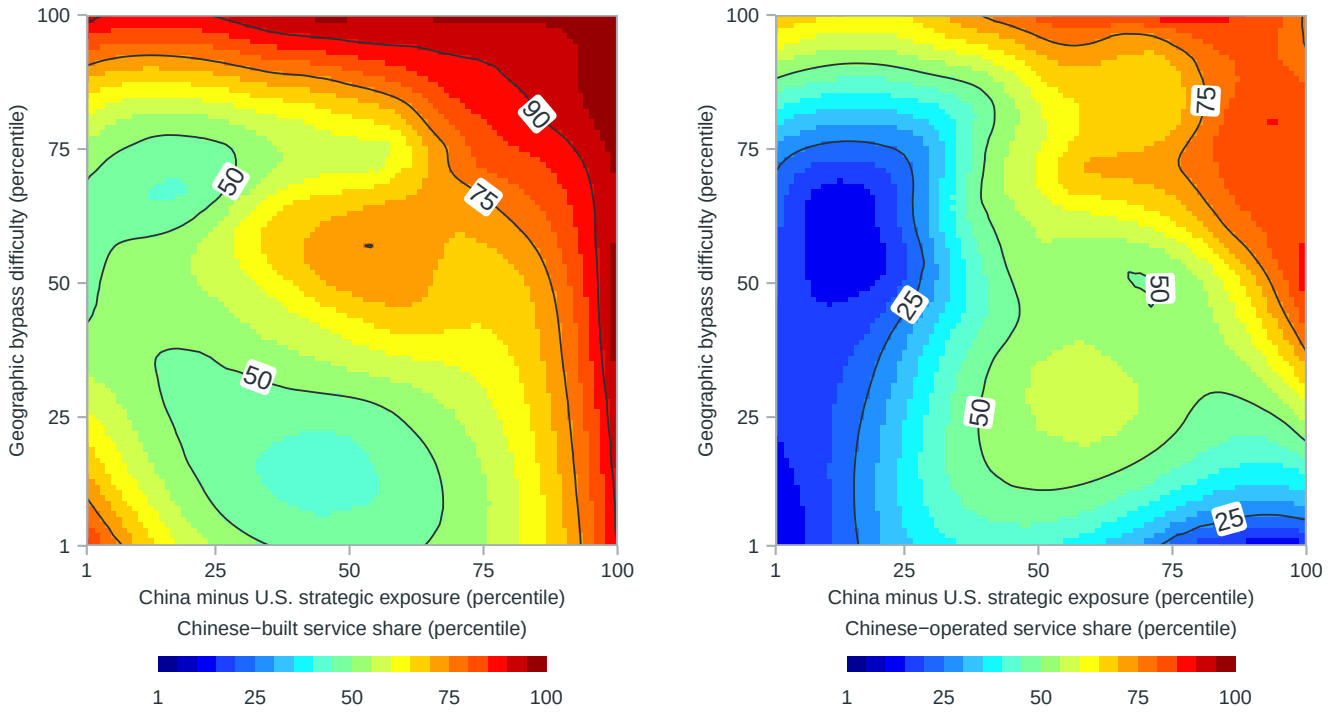
share of lane-level trade exposure is associated with 127 percent more capacity operated by its firms and 33 percent more capacity supplied by vessels built there. Both relationships remain positive when lanes beginning or ending in the country are excluded.

We also examine whether capacity is associated with trade exposure attributable to allies or non-allied countries. Allied trade exposure is positively associated with capacity on vessels built in the country, but not with capacity operated by its firms. For non-allied trade exposure, the association with capacity operated by the country's firms is marginally significant ($p = 0.09$), and the interactions with the difficulty of rerouting are not statistically significant. The association with capacity built in the country is small and positive, significant at the 10 percent level, and at the 5 percent level when the rerouting interactions are included.

Stylized Fact 4 (Operator and builder fleet vulnerability are distinct). Shipbuilding is geographically concentrated, but where a vessel was built does not identify who operates it: roughly three-quarters of Chinese-built capacity is operated by non-Chinese carriers. Dependence on vessels built in a country differs from dependence on vessels operated by its firms, so restrictions on each affect different parts of the network.

3.5 Combining Sources of Maritime Vulnerability

We conclude by bringing together trade exposure, geographic vulnerability, and dependence on foreign fleets to examine how much they overlap. Using the United States and China as an illustration, we ask whether the lanes that matter more for one country's trade are also harder to bypass and more dependent on Chinese-built or Chinese-operated vessels. This comparison highlights where these sources of maritime dependence coincide. The subsequent equilibrium analysis then determines what they imply



(a) Chinese-built service share

(b) Chinese-operated service share

Figure 9. Exposure, geography, and shipping capacity. The horizontal axis ranks the China-minus-U.S. relative trade exposure (labeled strategic exposure); the vertical axis ranks capacity-weighted detour severity (labeled geographic bypass difficulty). Color and contours report fitted Chinese builder-share percentiles in panel (a) and operator-share percentiles in panel (b), using bivariate LOESS. Red areas indicate higher fitted shares. Source: assigned trade, geographic disruptions, and vessel registries.

for shipping costs and welfare.

Figure 9 shows two heatmaps that bring the three dimensions together. The horizontal axis ranks lanes by relative trade exposure, the China-minus-U.S. gap, so lanes farther right matter more for Chinese trade. The vertical axis ranks lanes by geographic vulnerability, the average extra sailing distance needed to bypass local disruptions, so lanes higher up are harder to bypass. The colors and contours show the estimated percentile of dependence on Chinese vessels at each point, measured by the capacity share of Chinese-built vessels in panel (a) and Chinese-operated vessels in panel (b). Red regions mean greater dependence. The upper right is where these sources of maritime dependence coincide, with greater trade exposure for China, lanes that are harder to bypass, and greater reliance on Chinese vessels. Note that the distributions in both panels can differ since a lane can depend heavily on Chinese-built vessels without depending heavily on Chinese operators.

Stylized Fact 5 (Exposure and sources of vulnerability). The cost of a maritime disruption depends on the trade exposed, the availability of alternative routes, and the vessel capacity that can serve them. For the United States and China, the lanes that matter more for Chinese trade also tend to be harder to bypass and more reliant on Chinese-built vessels.

4 A Quantitative Model of Containerized Shipping

Section 2 developed an analytical framework for maritime vulnerability, and Section 3 measured its empirical components in the global shipping network. Here, we integrate the two in a spatial equilibrium model with endogenous routing and transport supply, where disruptions to maritime geography and fleet availability map into shipping costs, trade, and welfare. The model determines the disruption response $\phi_{e,g}$ and the producer-price responses $r_{o,g}$ of Section 2. The remainder of the paper takes this model to the data, examines its mechanisms, and conducts counterfactual exercises.

Our model features a trade block and a shipping block. The trade block takes generalized edge costs as given and determines routes, market access, and traffic on each directed edge, hence generating demand for shipping services. This demand is met by the shipping block, which allocates the available fleet across rotations and determines the supply of shipping services on each edge. Freight rates are equilibrium outcomes that adjust until demand equals supply on every edge. Since these rates enter into the generalized edge costs, routes, traffic, fleet deployment, and shipping costs are jointly determined in equilibrium.

Figure 10 summarizes the model. In the trade block (left), goods produced at each location are shipped to consumers elsewhere, and cargo is routed through the network, with cheaper paths attracting more traffic. Wages, output, and the location of workers within countries adjust to these shipping costs, and in turn shape how much cargo moves across each edge. In the shipping block (right), each operator pool allocates its fleet across vessel rotations. Serving an edge requires a complete round trip: capacity must reach the loading port, carry cargo, and return to its base port. A disruption anywhere along a rotation therefore reduces the services the fleet can supply, and because each pool serves many edges, its effects spread across the network. The two blocks meet in the market for shipping services (bottom): freight rates clear this market on each edge and, together with port congestion, determine the shipping costs that guide routing.

We begin by describing the economic environment. We then characterize the demand and supply of maritime transport and define the equilibrium.

4.1 Economic Environment

Geography and network. The world consists of M countries, indexed by $m = 1, \ldots, M$. Each country m contains a finite set of locations $\mathcal{N}_m$, and $\mathcal{N} = \bigcup_{m=1}^{M} \mathcal{N}_m$ denotes the set of locations globally. Each location is a port on the maritime transport network, and ports are connected by a directed graph $G = (\mathcal{N}, \mathcal{E})$, where edge $e = (i, j)$ allows shipments to move from port i to port j. The graph need not be symmetric: the presence of (i, j) does not imply the presence of (j, i). The map $m(i)$ identifies the country containing port i.

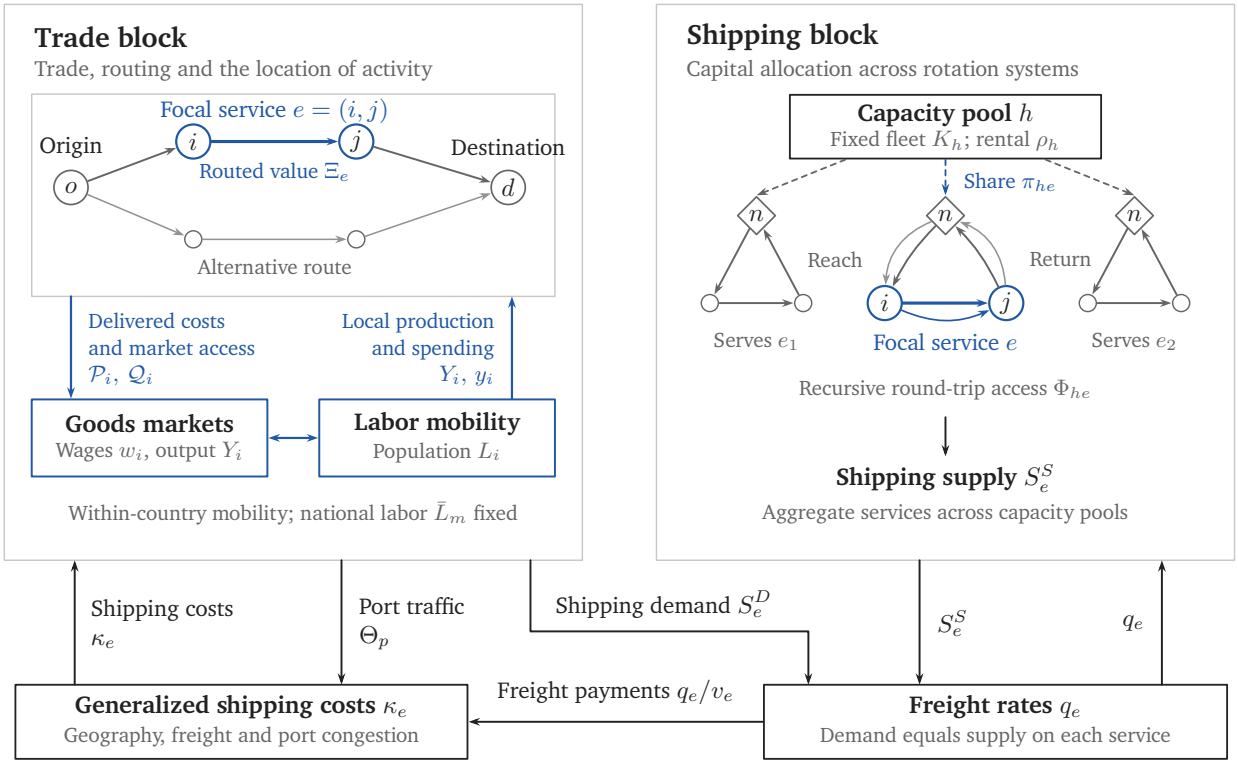


Figure 10. The joint equilibrium of the trade and shipping blocks. Left: the trade block, in which cargo from origin o to destination d is routed through the network, including service $e = (i, j)$. Right: the shipping block, in which a capacity pool allocates its fleet across rotations that start and end at a base port n. Bottom: freight rates clear the market for shipping services on each edge and, together with port congestion, determine generalized shipping costs. Solid arrows show cargo and vessel movements; dashed arrows show allocation decisions. Paths are illustrative.

Households and labor mobility. Each location j is populated by a mass L_j of workers who earn the local wage w_j. Workers value the local amenity u_j and consume varieties indexed by their location of origin. Preferences are

$$U_j = u_j c_j, \qquad c_j = \left[\sum_{i \in \mathcal{N}} c_{ij}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 1,$$

where c_{ij} denotes how much of location i's variety a worker at j consumes, and σ is the elasticity of substitution across varieties.

Local amenities may depend on population according to

$$u_j = \bar{u}_j L_j^{\beta}, \tag{4.1}$$

where $\bar{u}_j$ is the exogenous component of local amenities and β governs how amenities vary with local population. Workers take local amenities as given when choosing where to live.

The variety produced at i has factory-gate price p_i and delivered-price index $p_i \tau_{ij}$ at j. Here $\tau_{ij} > 0$

is the endogenous trade cost from i to j, which depends on all available routes and is derived below. A worker at j chooses $\{c_{ij}\}_{i\in\mathcal{N}}$ to maximize c_j subject to $\sum_i p_i\tau_{ij}c_{ij} \leq w_j$. The corresponding price index is

$$P_j = \left[\sum_{i\in\mathcal{N}} (p_i\tau_{ij})^{1-\sigma}\right]^{1/(1-\sigma)}.$$

Let $X_{ij} \equiv p_i\tau_{ij}L_jc_{ij}$ denote the total expenditure by workers at location j on the variety produced at i. Since workers spend all their labor income, total expenditure at j is $y_j \equiv w_jL_j = \sum_{i\in\mathcal{N}} X_{ij}$. CES demand implies that the income j spends on variety i depends on its relative delivered price $\left(\frac{p_i\tau_{ij}}{P_j}\right)$:

$$X_{ij} = \left(\frac{p_i\tau_{ij}}{P_j}\right)^{1-\sigma} y_j.$$

Workers may move freely among locations within their country but not across countries. Each country m has a fixed labor force $\bar{L}_m$. In equilibrium, free mobility equalizes indirect utility across all populated locations within a country. Since the indirect utility at i is $\frac{u_iw_i}{P_i}$, this equalization condition requires, for every $i \in \mathcal{N}_m$ with $L_i > 0$,

$$\frac{u_iw_i}{P_i} = W_m, \tag{4.2}$$

where W_m denotes the common welfare level of workers in country m. In addition, the population across locations must sum to the country's fixed labor force:

$$\sum_{i\in\mathcal{N}_m} L_i = \bar{L}_m. \tag{4.3}$$

Production and trade. Each location i produces a differentiated variety using labor as its only input:

$$Y_i = A_iL_i, \qquad A_i = \bar{A}_iL_i^{\alpha}, \tag{4.4}$$

where $\bar{A}_i$ is the exogenous component of local productivity and α governs how productivity varies with local employment. Firms take productivity A_i as given. Producing one unit requires $1/A_i$ units of labor, so perfect competition implies a factory-gate price $p_i = w_i/A_i$.

Substituting the competitive price into household demand yields the gravity-type expression for j's expenditure on variety i

$$X_{ij} = A_i^{\theta}w_i^{-\theta}\tau_{ij}^{-\theta}P_j^{\theta}y_j, \qquad \theta \equiv \sigma - 1 > 0. \tag{4.5}$$

Goods-market clearing requires the value of production at location i to equal total expenditure on its variety. Since $p_iY_i = w_iL_i = y_i$, this condition is

$$y_i = \sum_{j\in\mathcal{N}} X_{ij}. \tag{4.6}$$

4.2 Demand and Supply of Maritime Transport

4.2.1 Trade Block: Demand for Maritime Transport

Up to this point, the household problem pins down the bilateral expenditure X_{ij} for a given cost of delivering goods from i to j. However, it does not determine which routes carry those goods. Here we build trade costs from the costs of the available routes and split bilateral trade across them, thereby pinning down demand for shipping services on each edge.

There are three ingredients. First, a route is an ordered sequence of directed edges. Second, iceberg costs are multiplicative along that sequence. Third, routing is subject to idiosyncratic productivities that allow shipments between the same origin and destination to use different routes. We denote $N^+(i)$ and $N^-(i)$ as the outgoing and incoming neighbors of location i.

We start by considering a shipment at location i which is destined for location j. Sending it through an outgoing neighbor $k \in N^+(i)$ incurs the generalized edge cost κ_{ik} and leaves the continuation cost τ_{kj} from k to j. The deterministic cost of this continuation is therefore $\kappa_{ik}\tau_{kj}$. Routing selects the continuation with the lowest cost after accounting for an idiosyncratic routing productivity.

We assume that routing efficiencies are drawn independently from a Fréchet distribution across the feasible continuations, with a common shape parameter $\theta = \sigma - 1$. This assumption implies that lower-cost continuations carry more traffic, but multiple routes can be used at the same time. The parameter θ governs how responsive route shares are to relative costs, and is common across origins and routes (Allen, Fuchs and Wong, 2025; Fuchs and Wong, 2026). Since trade costs are aggregated over all the feasible routes from i to j, trade cost τ_{ij} is determined by the following:

$$\tau_{ij}^{-\theta} = \mathbf{1}\{i = j\} + \sum_{k \in N^+(i)} \left(\kappa_{ik}\tau_{kj}\right)^{-\theta}. \tag{4.7}$$

The first term in Equation (4.7) captures the local route that needs no further shipping, so it equals one if the current location i is the destination j and zero otherwise. The second term sums over the next stops k, each contributing the cost κ_{ik} of reaching k and the cost τ_{kj} of continuing from there to j. Routes may pass through a location more than once, so the sum includes routes of every finite length. Conditional on destination j, the share of shipments continuing through k is $(\kappa_{ik}\tau_{kj})^{-\theta}/\tau_{ij}^{-\theta}$.

The remaining problem is to aggregate these bilateral routing decisions into total traffic on each edge. Doing this directly would require tracking every origin–destination flow over every feasible route. Instead, the recursive structure here provides a simpler solution. Any route that uses edge $e = (i, j)$ can be split into three parts: from its origin to i, along edge e, and from j to its destination. We can then add up each part separately.

Let inward market access $\mathcal{P}_i$ measure the supply that can reach i, the start of an edge. It consists of

local supply and the supply arriving from upstream locations:

$$\mathcal{P}_i = A_i^{\theta} w_i^{-\theta} + \sum_{k \in N^-(i)} \kappa_{ki}^{-\theta} \mathcal{P}_k. \tag{4.8}$$

Similarly, let outward market access $\mathcal{Q}_i$ measure the expenditure that can be reached from i, the end of an edge. It consists of local expenditure and the expenditure reachable through downstream locations:

$$\mathcal{Q}_i = P_i^{\theta} y_i + \sum_{k \in N^+(i)} \kappa_{ik}^{-\theta} \mathcal{Q}_k. \tag{4.9}$$

We now put together upstream supply, the cost of crossing the edge, and downstream expenditure. Using Equations (4.8) and (4.9), total value traffic on edge $e = (i, j)$ is

$$\Xi_e = \kappa_e^{-\theta} \mathcal{P}_i \mathcal{Q}_j. \tag{4.10}$$

Equation (4.10), derived in Appendix A.3, is the network version of a gravity equation, with edges in place of country pairs. Traffic on an edge is high when many sellers can reach its start (supply), many buyers can be reached from its end (expenditure), and the edge is cheap to cross.

While the traffic measure Ξ_e in equation (4.10) is expressed in value terms, the shipping block operates in physical units of transport. To link the two blocks, we introduce $v_e > 0$, an exogenous edge-specific parameter that measures the amount of cargo value per unit of shipping services. Demand for shipping services on edge e is therefore

$$S_e^D = \frac{\Xi_e}{v_e}, \tag{4.11}$$

where a lower v_e means that a given value of trade requires more shipping services.

4.2.2 Shipping Block: Supply of Maritime Transport

The previous subsection determined demand for shipping services on each edge. We now describe who supplies these services. The shipping block takes freight rates as given and determines capacity rentals, vessel deployment, and the supply of shipping services on each edge. We start by introducing the shipping operators and ship rotation technology. We then outline the operator's deployment problem and show how maritime geography determines productivity across rotations. Finally, we derive capacity rentals, deployment shares, and edge-level supply.

Shipping operators. The shipping sector consists of competitive operators. Since operators decide where to deploy their vessels, we group vessel capacity into pools h, each with a fixed stock K_h. Capacity can be redeployed within a pool, but not between pools. In the quantitative analysis, we assign the

vessels run by operators from one country to the same pool. Renting capacity in pool h costs ρ_h, and shipping services on edge e sell at the freight rate q_e. Operators take both prices as given.

Rotation technology. Like buses on a fixed route schedule, container ships follow circular sequences of port calls (rotations). Our rotation data in Section 3 record these sequences, and consecutive port calls define the directed edges on which a vessel provides shipping capacity. Our model captures this circular structure by requiring vessel capacity to operate on complete rotations. Each rotation starts and ends at a base port n, which can differ from the port where cargo is loaded. The weights $b_{hn} \geq 0$, with $\sum_n b_{hn} = 1$, capture the share of pool h's rotations based at port n. To supply shipping services on edge $e = (i, j)$, capacity from pool h and base port n must reach loading port i, carry cargo from i to j, and return to n. Each part of the rotation may pass through intermediate ports, allowing its physical path to reflect the circular, multi-port itineraries observed in the data. The right-hand block of Figure 10 illustrates the case of a single base port.

The requirement to return is what distinguishes containership vessels from cargo: a shipment can stop at its destination, whereas a vessel that delivers its cargo at j must still sail back to its base port. Each unit of capacity serves a single rotation. In the implementation of our quantitative analysis, observed vessel rotations define the network and the baseline allocation of capacity across edges. In counterfactuals, however, capacity is not tied to the observed schedules: it can be reassigned to any feasible rotation through the network.

Operator's problem. A unit of capacity from pool h can be deployed on any eligible rotation serving an edge e. Let Z_{he} denote the units of shipping services produced by that capacity when assigned to such a rotation. At freight rate q_e, the deployment generates revenue $q_e Z_{he}$. After renting the unit of capacity at rate ρ_h, the operator solves

$$\max_{e \in \mathcal{E}} \{q_e Z_{he} - \rho_h\}.$$

Because ρ_h is common to all uses of capacity in pool h, the operator assigns the unit to the rotation with the highest shipping revenue $q_e Z_{he}$.

Within a pool, some capacity fits certain rotations better than others. We model this heterogeneity with independent Fréchet draws Z_{he} across eligible edges:

$$\Pr(Z_{he} \leq z) = \exp\{-\Phi_{he} z^{-\zeta}\}, \qquad \zeta > 1.$$

The parameter Φ_{he} is the scale of this productivity distribution. A higher Φ_{he} makes high values of Z_{he} more likely. Defining $\gamma_\zeta \equiv \Gamma(1 - 1/\zeta)$, the expected productivity is

$$\mathbb{E}[Z_{he}] = \gamma_\zeta \Phi_{he}^{1/\zeta},$$

where $\Phi_{he}^{1/\zeta}$ is proportional to the shipping services that one unit of capacity from pool h is expected to produce on edge e. We now construct Φ_{he} from the maritime network.

Round-trip access. Conditional on base port n, a rotation serving edge $e = (i, j)$ has three parts (indexed by $X \in \{B, L, R\}$): the ballast trip from n to loading port i (B), the laden cargo trip from i to j (L), and the return to n (R). Each part may pass through several intermediate ports and can therefore be completed along many paths, which reflects the indirect nature of container-shipping rotations (Ganapati, Wong and Ziv, 2024). We call Φ_{he} the round-trip access, the theoretical counterpart of the round-trip effect in Wong (2022), because it depends on the complete rotation needed to serve the edge.

For each part $X \in \{B, L, R\}$ of a rotation, let $g_{ik}^X > 0$ be the operating requirement of moving vessel capacity from location i to a feasible next location k, such as sea travel distance and time. Operators take this as given. In our geographic counterfactuals this parameter will increase if a disruption makes the trip longer or more demanding. We calibrate the baseline values to observed traffic, and disruptions scale them by the resulting increase in sailing distance (Appendix B.3).

Directly enumerating every feasible walk is computationally intractable because the network contains cycles and arbitrarily long paths. Instead, we calculate the access terms recursively for each rotation part $X \in \{B, L, R\}$:

$$\mathcal{R}_{ij}^X = \mathbf{1}\{i = j\} + \sum_k (g_{ik}^X)^{-\zeta} \mathcal{R}_{kj}^X. \tag{4.12}$$

The first term in Equation (4.12) is the local route that needs no further sailing, and the remaining terms add every feasible continuation, including routes that return. Access is higher when the operating requirement g_{ik}^X falls, and when more routes are available (adding more terms to the sum over k).

The round-trip access term then combines the three parts of a rotation mentioned above:

$$\Phi_{he} = a_{he}^{-\zeta} \sum_n b_{hn} \underbrace{\mathcal{R}_{n,i(e)}^B}_{\text{Ballast}} \underbrace{\mathcal{R}_{i(e),j(e)}^L}_{\text{Laden}} \underbrace{\mathcal{R}_{j(e),n}^R}_{\text{Return}}. \tag{4.13}$$

Each $\mathcal{R}$ term is the access for one part of the rotation: reaching the loading port (B), serving the laden leg (L), and returning to the base port (R). It is higher when any one part is shorter or less demanding and when more routes can complete it. The three terms multiply, so a more difficult or longer trip on any one part decreases Φ_{he}. The sum over n averages across the base ports of pool h, using the weights b_{hn}. The wedge $a_{he} \geq 1$ raises the operating requirement of pool h on edge e, for example through an operational restriction, and Φ_{he} falls as it rises. Overall, Φ_{he} governs how productive a pool is on a rotation, while K_h governs how much capacity it has.

Of the three parts, only the laden part of a rotation earns revenue. The ballast and return parts earn nothing for edge e. A real vessel may still carry cargo on those legs, but the model treats that as a rotation serving a different edge. A closure that lengthens the return from j, for example, reduces the

services a pool can supply on edge e even if the edge itself is unaffected. An operational restriction works differently: it raises a_{he} only for the targeted pools and edges, so the route stays open to other capacity.

Capacity allocation. Because operators are competitive, the rental rate ρ_h equals the average revenue that units of capacity in pool h earn on their best rotations. With Fréchet-distributed productivity, this average has a closed form:

$$\rho_h = \gamma_\zeta \left(\sum_{e \in \mathcal{E}} q_e^\zeta \Phi_{he} \right)^{1/\zeta}. \tag{4.14}$$

The rental rate summarizes the value of all rotations available to capacity in pool h. A valuable shipping opportunity on any edge raises the opportunity cost of deploying that capacity elsewhere.

As a result, the share of capacity in pool h assigned to rotations serving edge e is

$$\pi_{he} = \frac{q_e^\zeta \Phi_{he}}{\sum_{e' \in \mathcal{E}} q_{e'}^\zeta \Phi_{he'}}.$$

Capacity is drawn toward edges with higher freight rates and greater round-trip access. The parameter ζ governs how strongly deployment responds to differences in freight rates across edges.

Supply of shipping services. We now aggregate operators' deployment decisions into the supply of shipping services on each edge. The rental rate ρ_h is the revenue per unit of capacity in pool h, with each unit deployed on its best rotation. A marginal increase in the freight rate q_e raises this revenue in proportion to the shipping services the pool supplies on edge e. By the envelope theorem, $\partial \rho_h / \partial q_e$ is therefore the quantity of shipping services supplied on edge e per unit of capacity in pool h.

Multiplying by the available capacity K_h and summing across pools gives

$$S_e^S = \sum_h K_h \frac{\partial \rho_h}{\partial q_e} = \gamma_\zeta^\zeta q_e^{\zeta - 1} \sum_h \Phi_{he} \rho_h^{1-\zeta} K_h. \tag{4.15}$$

Supply on edge e is greater when its freight rate is high, when more capacity can serve the required rotation, and when that capacity has less valuable alternatives elsewhere. Because capacity in the same pool can serve several edges, a higher freight rate on one edge can draw capacity away from other rotations. A disruption that lowers access on some rotations, or restricts some pools, can therefore move freight rates on every edge those pools serve.

Since operators are competitive, the shipping block equates total service revenue to capacity rentals:

$$\sum_e q_e S_e^S = \sum_h \rho_h K_h.$$

This formulation combines the Eaton–Kortum allocation mechanism with the complete-rotation technology of maritime transport (Eaton and Kortum, 2002; Wong, 2022; Kalouptsidi, 2014; Brancaccio,

Kalouptsidi and Papageorgiou, 2020; Ganapati and Wong, 2023; Dunn and Leibovici, 2024). Appendix B.3 provides the derivations and regularity conditions.

Disruptions and freight rates. A disruption affects the supply of shipping services on an edge in two ways. Directly, it lowers the round-trip access of the pools that serve the edge, either because a closure lengthens part of their rotations or because a restriction targets particular pools. Indirectly, it changes what those pools can earn elsewhere. When freight rates rise on some edges, capacity moves toward them and away from others. Freight rates then adjust until supply equals demand on every edge, and these changes enter the edge-cost responses $\phi_{e,g}$. The rotation elasticity ζ governs the indirect channel. When ζ is close to one, capacity barely moves across edges, so freight rates absorb changes in demand edge by edge. When ζ is higher, capacity moves readily to wherever it earns the most, so freight rates change in similar proportions across edges served by the same pools. Because rotations are round trips, a closure lowers access on every edge whose rotation crosses it, not only on edges whose cargo crosses it.

4.2.3 Port Congestion and Generalized Edge Costs

The demand and shipping-supply blocks determine traffic and freight rates. Traffic may also affect transport costs by congesting ports with limited handling capacity. We model this channel separately from shipping operators' vessel-deployment problem. Generalized edge costs combine port congestion with freight rates and the exogenous characteristics of each edge.

Port congestion. Cargo must be handled at the ports where each edge begins and ends, and ports become costlier to use as they get busier. Let $\mathcal{E}_p$ denote the edges that begin or end at port p, K_p^P its handling capacity, and $\iota_{pe} \geq 0$ the port capacity used at p per unit of traffic on edge e. Port throughput and the resulting cost factor are

$$\Theta_p = \sum_{e \in \mathcal{E}_p} \iota_{pe} \Xi_e, \qquad D_p = \left(\frac{\Theta_p}{K_p^P} \right)^{\lambda_p}, \tag{4.16}$$

where $\lambda_p \geq 0$ governs the strength of congestion. When $\lambda_p > 0$, more traffic relative to capacity raises the cost of every edge using the port. If $\lambda_p = 0$, this congestion channel shuts down. Throughput counts every container handled, so even transshipped cargo adds to congestion at each port at which it calls.

Generalized edge costs. The transport-demand block takes the generalized edge cost κ_e as given. We can now determine this cost from the exogenous characteristics of the edge, congestion at its origin and destination ports, and the freight rate. For an edge $e = (i, j)$, let $\bar{\kappa}_e$ denote the exogenous component of its cost, including its geographic characteristics and any primitive cost wedge. The generalized edge cost

can be then written as follows:

$$\kappa_e = \underbrace{\bar{\kappa}_e}_{\text{exogenous cost}} \times \underbrace{D_i D_j}_{\text{port congestion}} \times \underbrace{\left(1 + \frac{q_e}{v_e}\right)}_{\text{freight}}, \tag{4.17}$$

where D_i and D_j are the congestion cost factors at the origin and destination ports, and q_e/v_e is the freight payment per unit of cargo value. Equation (4.17) combines geography, congestion, and the endogenous freight rate. This is the mapping to edge costs of Section 2 in our model. Differentiating it with respect to the intensity d_g of scenario g gives the edge-cost response $\phi_{e,g}$,

$$\phi_{e,g} = \underbrace{\frac{\partial \log \bar{\kappa}_e}{\partial d_g}}_{\text{exogenous cost}} + \underbrace{\sum_{p \in \{i,j\}} \lambda_p \left(\frac{\partial \log \Theta_p}{\partial d_g} - \frac{\partial \log K_p^P}{\partial d_g} \right)}_{\text{port congestion}} + \xi_e \underbrace{\left(\frac{\partial \log q_e}{\partial d_g} - \frac{\partial \log v_e}{\partial d_g} \right)}_{\text{freight response}}, \tag{4.18}$$

where the second term sums over the two ports at the ends of edge $e = (i, j)$, and $\xi_e \equiv (q_e/v_e)/(1+q_e/v_e)$ is the baseline share of the freight rate in the generalized edge cost. The first term changes only on the edges a disruption affects directly; for closures, it reflects the detours measured in Section 3. Capacity withdrawals and port fees act through the freight term. The congestion and freight terms can change on any edge, because ports and the fleet are shared.

4.3 Competitive Equilibrium

We now define the competitive equilibrium.

Definition 4.1 (Competitive equilibrium)**.** Given the network and primitives, a competitive equilibrium consists of prices and allocations satisfying the following conditions:

(i) *Households.* Given wages, delivered goods prices, and local amenities, households choose consumption and location to maximize utility.

(ii) *Goods producers.* Given wages and goods prices, producers choose employment and output to maximize profits.

(iii) *Shipping operators.* Given freight rates and capacity rentals, shipping operators rent vessel capacity and deploy it across eligible rotations to maximize profits. Competition for the fixed capacity K_h determines its rental rate ρ_h, and operators allocate that capacity according to the deployment shares π_{he}.

(iv) *Trade routing and transport costs.* Bilateral shipments follow the route-choice system, while port throughput and generalized edge costs are consistent with traffic and freight rates.

(v) *Labor markets.* The allocation of workers across locations exhausts each country's labor supply:

$$\sum_{i \in \mathcal{N}_m} L_i = \bar{L}_m, \qquad m = 1, \ldots, M. \tag{4.19}$$

(vi) *Goods markets.* The value of production at each location equals total expenditure on its variety:

$$p_i Y_i = w_i L_i = \sum_{j \in \mathcal{N}} X_{ij}, \qquad i \in \mathcal{N}. \tag{4.20}$$

(vii) *Markets for shipping services.* On each edge, the demand for shipping services equals their supply:

$$S_e^D = S_e^S, \qquad e \in \mathcal{E}. \tag{4.21}$$

Wages and goods prices clear labor and goods markets across locations. Freight rates clear markets for shipping services edge by edge, while capacity rentals price vessel capacity across competing rotations. Because port costs depend on total traffic, route choice, freight rates, and vessel deployment are jointly determined in equilibrium.

4.4 Counterfactuals

We evaluate counterfactuals as changes relative to the observed baseline equilibrium. Following the exact-hat algebra commonly used in quantitative trade models, we express each model object relative to its baseline value:

$$\widehat{x} = \frac{x'}{x^0},$$

where x^0 and x' are positive baseline and counterfactual values. The hat system expresses the recursive equilibrium in proportional changes. For a baseline satisfying the level equations, exact hats remove many unobserved levels (Dekle, Eaton and Kortum, 2008; Costinot and Rodríguez-Clare, 2014). The counterfactual system requires only three sets of baseline shares: (1) routing shares, which split each location's throughput between local production or spending and traffic on each edge; (2) revenue shares, which give each pool's freight revenue by edge and each edge's freight revenue by pool; and (3) cost shares, which give the freight share of each edge's cost and each edge's share of port throughput.

Counterfactual changes. A counterfactual consists of changes to one or more exogenous model objects. A geographic counterfactual changes the exogenous cost of traversing an edge, $\widehat{\bar{\kappa}}_e$, and the round-trip access of capacity based at h to edge e, $\widehat{\Phi}_{he}$ (Section 6.1). The former captures changes in the geographic or policy cost of using the edge, while the latter captures changes in the requirements for reaching the loading port, carrying cargo across the edge, and returning to base. A fleet withdrawal raises a_{he} for the targeted capacity on the targeted lanes, without changing K_h (Section 7.1). A port fee drives a

wedge between the freight rate that shippers pay and what the carrying pool receives (Section 7.2). Some disruptions apply only to certain cargo, depending on its origin or destination (Section 6.2). For these, cargo subject to the restriction faces a different network than other cargo, while all cargo shares the same vessels. Conditional on these changes, prices, quantities, routes, and capacity allocations are determined endogenously.

Proposition B.1 in Appendix B.2 states the counterfactual equilibrium in proportional changes. This representation takes the directed network and the rotations available to each pool as given, so trade and vessel capacity reallocate among existing options. When a disruption leaves an edge without a route, we remove the edge and solve the equilibrium on the new network.

Welfare costs. We measure country m's welfare cost of counterfactual g as

$$C_m(g) = -100 \log \widehat{W}_m(g), \tag{4.22}$$

in log percentage points relative to the baseline equilibrium unless stated otherwise, where W_m is the common utility level of workers in that country. Since utility is linear in the consumption aggregate, this measure has a consumption-equivalent interpretation: a value of one means that counterfactual consumption would need to increase by approximately 1 percent to restore baseline utility. Positive values therefore denote welfare losses, while negative values denote gains. To first order, $C_m(g)$ equals the equivalent-variation loss of Proposition 2.1 for a disruption of intensity d_g, expressed as a percent of baseline income.

4.5 Taking the Model to the Data

The quantitative implementation maps observed ports and shipping lanes into the locations and edges of the model and constructs baseline measures of trade, route traffic, vessel capacity, and its allocation across the network. We combine the data described in Section 3 with national accounts to assign income, labor, and expenditure across locations and obtain the baseline equilibrium.

Geography. The model network is built from the schedules described in Section 3.1 and covers 77 economies. Within these economies, 196 major ports are separate locations, and each economy's remaining ports, if any, are combined into one location, giving 245 locations in all. The network has 3,402 directed edges. The lane experiments in Section 5.1 use 1,031 international edges, each linking two major ports.

Vessel operators and builders. The data record each vessel's operator country and builder country. Because operators control deployment, we define pools by operator country, so h indexes operator

countries. Each pool's capacity is fixed: it can be reallocated across the edges its operators serve in the baseline, but not across pools.[5]

Shipbuilders are not decision makers in the model. Builder country instead describes the composition of the capacity controlled by each operator country. We retain this information to identify builder-linked exposure. An operator-linked counterfactual applies to capacity controlled by operators from the targeted country, whereas a builder-linked counterfactual applies to vessels built in the targeted country, regardless of who operates them.

National accounts. Country GDP and population determine total baseline income and labor in each country. Because comparable measures are not available at the port level, we allocate these totals across port locations in proportion to each location's share of national incoming and outgoing vessel traffic. This allocation preserves each country's total income and labor.

4.6 Parameterization

We now complete the quantitative specification by assigning values to the model's elasticities and cost shares. These parameters govern substitution across varieties, local productivity and amenity effects, the allocation of vessel capacity across rotations, and the contribution of freight rates and port congestion to generalized edge costs. Table 1 summarizes the benchmark parameterization.

Trade and location. We set the elasticity of substitution across origin varieties to $\sigma = 4$, which implies a trade elasticity of $\theta = \sigma - 1 = 3$. The implied trade elasticity lies within the range commonly used in quantitative trade models (Simonovska and Waugh, 2014; Head and Mayer, 2014). Following the spatial specification in Section 4.1, productivity rises with local employment with elasticity $\alpha = 0.10$, while amenities fall with local population with elasticity $\beta = -0.30$, as in Allen and Arkolakis (2014). These two parameters allow economic activity and congestion to respond within countries rather than fixing the spatial distribution of activity.

Maritime transport. We set the rotation elasticity to $\zeta = 4$ and the baseline freight share to $\xi_e = 0.041$ on every edge. Thus a 1 percent increase in the freight rate raises generalized edge cost by approximately 0.041 percent, holding other components fixed. We measure this share in U.S. Census import data: in July–October 2023, freight charges net of insurance were 4.3 percent of the customs value of containerized imports. The global counterfactuals set port congestion to $\lambda_p = 0.096$.

[5]Vessels without a recorded operator country form a residual pool with 2.6 percent of service revenue.

Table 1. Parameterization

Parameter	Value	Role	Source
A. Trade and location			
Substitution elasticity σ	4	Substitution across origin varieties; implies $\theta = 3$.	Literature (Simonovska and Waugh, 2014; Head and Mayer, 2014)
Productivity elasticity α	0.10	Response of productivity to local employment.	Allen and Arkolakis (2014)
Amenity elasticity β	-0.30	Response of amenities to local population.	Allen and Arkolakis (2014)
B. Maritime transport			
Rotation elasticity ζ	4	Heterogeneity and capacity reallocation across rotations.	Set; not identified by the Red Sea event
Freight-rate share ξ_e	0.041	Pass-through of freight rates into generalized edge costs.	U.S. Census imports, July–October 2023
Port-congestion elasticity λ_p	0.096	Response of port costs to traffic relative to handling capacity.	Fuchs and Wong (2026); 0 in the Red Sea validation

4.7 Computational Approach

We compute counterfactual equilibria using the proportional-change system developed in Section 4.4. For each counterfactual, we specify changes in the relevant exogenous model objects and jointly solve for the resulting changes in trade, route traffic, freight rates, vessel deployment, wages, population, and welfare.

Computationally, we express the equilibrium conditions as a nonlinear system in log changes and solve it using a damped Newton-type method that starts at the baseline equilibrium. Each iteration uses the proposed changes in prices and market access to compute trade, traffic, and vessel deployment, and we stop when all equilibrium and market-clearing conditions hold. Appendices B.2 to B.3 present the system.

5 Model Diagnostics and External Validation

Before turning to the counterfactuals, we evaluate the model's implications along two dimensions. We first raise the cost of each shipping lane by the same amount, one at a time, to identify the lanes that matter most for each country's welfare and how far our measured trade exposure from Section 3.2 explains them. Second, we validate the model against the Red Sea crisis that began in late 2023 by comparing its predictions with the observed changes in routes, vessel deployment, freight charges, and U.S. imports.

5.1 Lane-Level Disruptions

We begin by investigating which shipping lanes matter most for welfare. For each international lane e, we increase its sailing requirement g_e by 10 percent, holding everything else fixed, and recompute the equilibrium. This lowers the round-trip access Φ_{he} of every rotation through the lane and raises the exogenous cost of shipping cargo on it by half that amount in logs, $\widehat{\kappa}_e = 1.10^{\frac{1}{2}}$. Because the shock

is identical across lanes, differences in welfare costs across lanes reflect countries' exposure and the equilibrium adjustment rather than the severity of the disruption.

This standardized exercise provides a benchmark for the counterfactuals that follow. Section 6 allows transport-cost changes to depend on maritime geography and the availability of alternative routes. Section 7 instead studies disruptions to the supply of vessel capacity. In both cases, the scenario-specific disruption interacts with the country exposure and equilibrium adjustment isolated here.

Baseline traffic is not sufficient to determine a lane's economic importance. A heavily used lane may have readily available substitutes, while a lane carrying less traffic may be costly to disrupt when alternative routes are limited. The welfare effect of disrupting a lane therefore depends on both country exposure and the scope for substitution in equilibrium.

The equilibrium response also means that the increase in the sailing requirement does not pass through one for one to the generalized cost faced by trade. Across the experiments, traffic on the disrupted lane falls by 13.5 log points on average, its freight rate rises by 5.5 log points, and its generalized cost rises by 4.9 log points. The freight rate rises even as traffic falls because the longer trip reduces the supply of shipping services on the lane. Port congestion also responds to traffic, providing an additional margin through which generalized edge costs adjust.

Measured exposure and welfare. We compare the China-minus-U.S. difference in welfare costs with the corresponding difference in strategic exposure (the relative trade exposure of Section 3.2), measured as each lane's share of the country's routed trade value.

Figure 11 documents a positive relationship between strategic exposure and the welfare-cost difference. A one-percentage-point increase in the China-minus-U.S. exposure difference is associated with a 0.0099 log-percentage-point increase in the corresponding welfare-cost difference. Strategic exposure accounts for 36.4 percent of the cross-lane variation in the welfare-cost difference. The fit comes from lanes that begin or end in China or the United States. On lanes between third countries, strategic exposure accounts for 10.3 percent of the variation, and on 6 percent of them both countries gain, through the responses of freight rates and producer prices.

The comparison distinguishes measured exposure from equilibrium welfare costs. Welfare costs also depend on the reallocation of vessel capacity and resulting freight-rate adjustments, port congestion, and the transmission of transport costs through equilibrium wages and prices. These adjustment margins account for the variation in welfare costs not summarized by measured exposure. Adding baseline traffic to strategic exposure raises the explained share of the variation across lanes from 36.4 to 74.1 percent, but the two together still leave much of it unexplained.

Decomposing welfare costs. Proposition 2.1 splits a country's welfare loss into the direct cost of its exposed trade and the adjustment of producer prices. Figure 12 shows that both matter. The direct cost accounts for 42.9 percent of the cross-lane variation in the China-minus-U.S. welfare-cost difference.

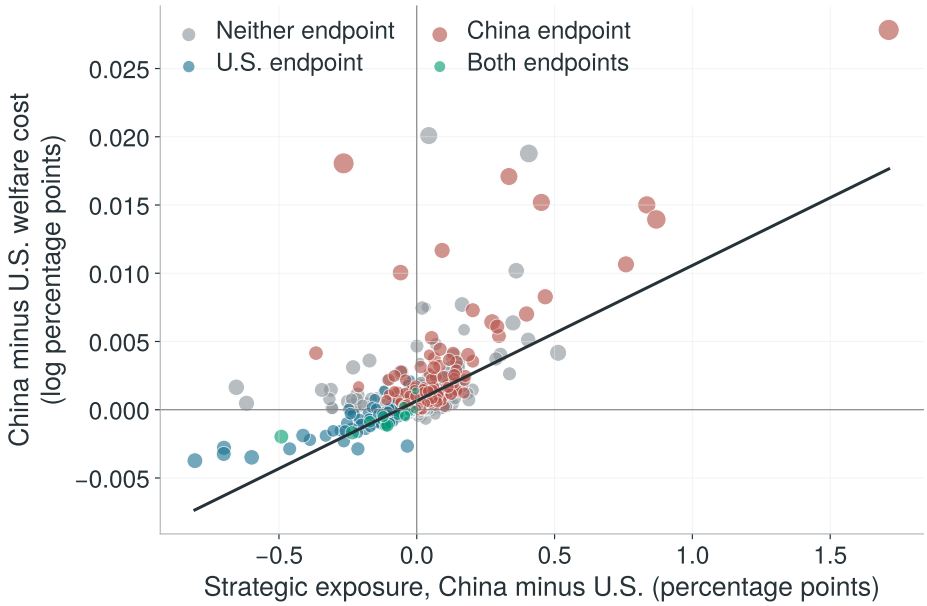


Figure 11. Strategic exposure and welfare costs. Each point is an international lane subject to a 10 percent increase in its sailing requirement. The horizontal axis reports the China-minus-U.S. difference in strategic exposure, in percentage points; the vertical axis reports the corresponding difference in welfare costs, in log percentage points. Bubble size increases with baseline traffic, and colors mark lanes with a Chinese, U.S., both, or neither endpoint. The line is an OLS fit, with slope 0.0099 and R^2 of 36.4 percent.

Adding producer-price adjustment raises this share to 99.8 percent. Price adjustment matters most on lanes that carry Chinese exports: a costlier lane lowers foreign demand for Chinese goods and, with it, the prices Chinese producers receive. Spillovers through the shipping network also matter. If we account only for the cost increase on the disrupted lane itself, ignoring how freight rates and port congestion spread it to other lanes, the share falls to 77.9 percent. Measured exposure thus captures the direct effect of a disruption, but its full welfare cost depends on how prices and the shipping network adjust.

5.2 Red Sea Disruptions

Beginning in late 2023, a series of Houthi attacks on commercial vessels in the Red Sea led major container carriers to avoid the region and reroute services that would ordinarily transit Suez around the Cape of Good Hope. We use this episode to assess the model's response to a geographic disruption before turning to the counterfactuals that follow. Conditional on the observed withdrawal from Suez, we investigate the extent to which the model accounts for changes in shipping routes, vessel deployment, freight charges, and import quantities.

To measure the shipping response, we construct capacity-weighted measures of passage and vessel deployment across routes and services. Dated carrier announcements (A.P. Moller–Maersk, 2023; Mediterranean Shipping Company, 2023; CMA CGM, 2023; Hapag-Lloyd, 2023) identify when each carrier suspended passage through the Red Sea. Vessel movements from Global Fishing Watch (GFW) (Global Fishing Watch, 2025) show whether ships passed through Suez, the Cape, or Panama before and after

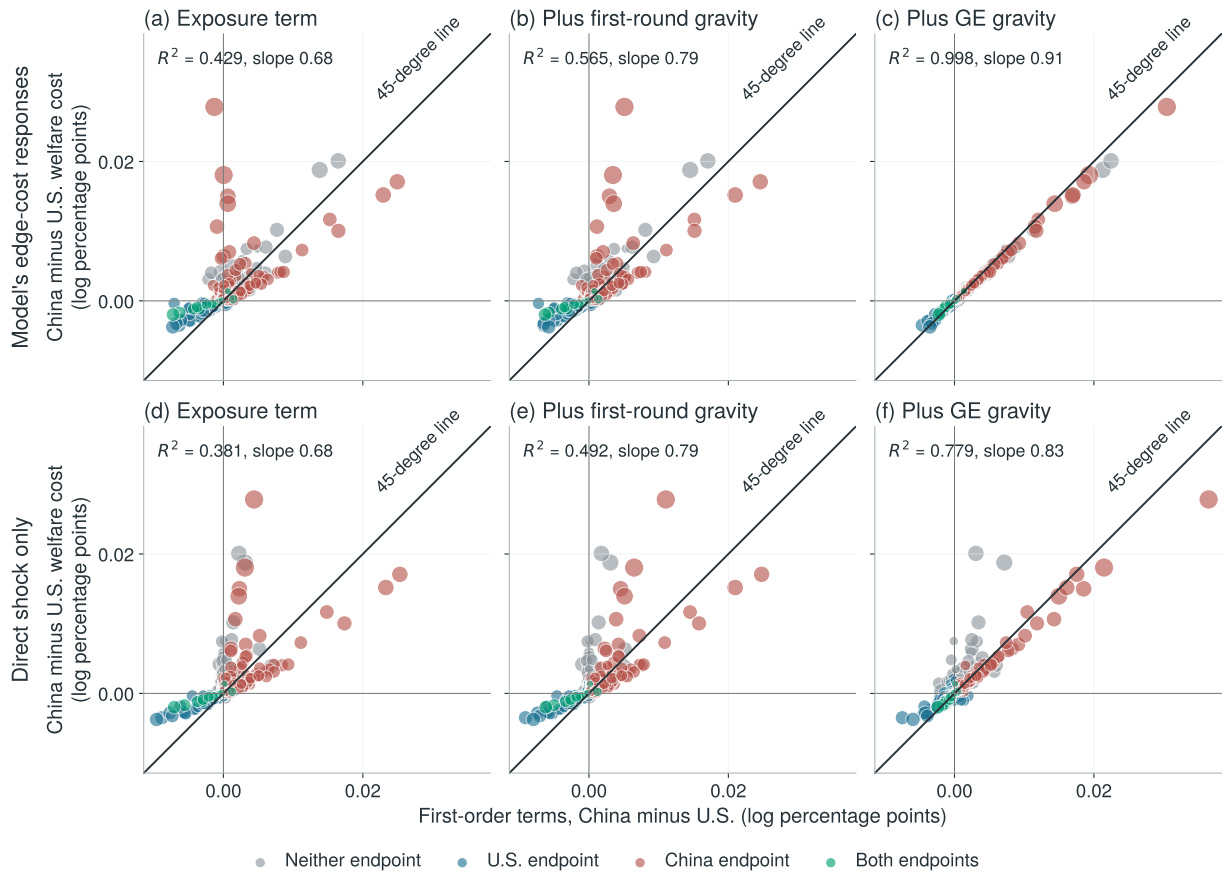
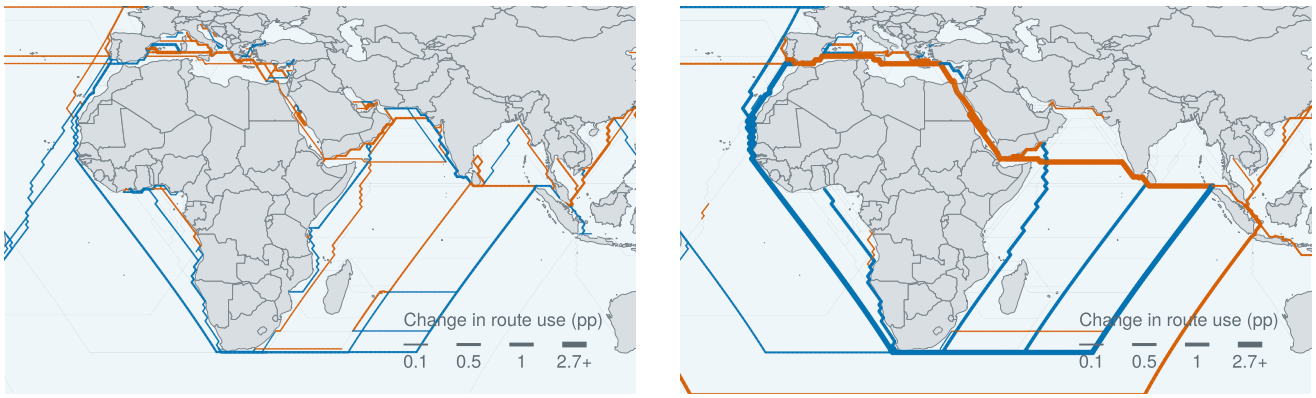


Figure 12. Exposure, prices, and network spillovers. Each point is a lane from Figure 11. The vertical axis reports the China-minus-U.S. welfare cost of a 10 percent increase in the lane's sailing requirement in the full model; the horizontal axis reports the prediction of Proposition 2.1. Panels (a) and (d) include only the direct cost of exposed trade; panels (b)–(c) and (e)–(f) add producer-price adjustment, with producer prices from partial and full gravity adjustment, respectively. The top row uses the model's edge-cost responses; the bottom row uses only the cost increase on the disrupted lane. Each panel reports the R^2 and slope of an OLS fit; the line is the 45-degree line. Both axes are in log percentage points.

the disruption, and how their commercial deployment changed. Vessel registries provide each ship's capacity, and scheduled rotations define the port-to-port service legs.

To measure the corresponding changes in U.S.-bound trade, we combine three additional data sources. Panjiva (S&P Global Market Intelligence, 2023–2024) records each U.S. import shipment's origin country, last foreign port, U.S. port of arrival, vessel, date, product, and volume; matching each vessel to its GFW track shows whether the cargo traveled through Suez, around the Cape of Good Hope, or through Panama. U.S. Census data (Census Bureau, 2023–2024) provide import quantities and freight charges by origin country, product, and customs district. PortWatch (Arslanalp et al., 2025) measures vessel transits through Suez, Bab el-Mandeb, the Cape, and Panama. Each source corresponds to a distinct model outcome: GFW to vessel routes and deployment, Panjiva to cargo routing, Census to import quantities and freight charges, and PortWatch to chokepoint passage. We compare each observed margin with its model counterpart separately.

Overall, between July–October 2023 and January–August 2024, the cargo carried by container ships



(a) Observed vessel reallocation **(b)** Model-implied vessel reallocation

Route use rises Route use falls

Figure 13. Vessel reallocation following the Red Sea disruptions. Panel (a) reports observed changes in route use; panel (b) reports the corresponding model-implied changes. Source: Global Fishing Watch vessel movements (Global Fishing Watch, 2025), vessel registries, and model estimates.

through the Suez Canal fell by 82 percent, from 34.6 million to 6.2 million metric tons a month, while the cargo carried by container ships around the Cape of Good Hope rose from 6.3 million to 35.1 million metric tons a month (Arslanalp et al., 2025). At the route level, Figure 13a shows how the capacity of the three carriers we track in GFW (Maersk, MSC, and Hapag-Lloyd) was reallocated across routes.

Model experiment. We use the Red Sea episode to evaluate the implications of the model. We model the episode as carrier-specific shocks that lower round-trip access Φ_{he} for operator-country pool h on each edge e whose sea path passes through Suez or Bab el-Mandeb. Carriers cut back their use of Suez to very different degrees, reflecting differences in their exposure to the attacks, insurance costs, and the profitability of their networks. We discipline these shocks with the observed withdrawals: the shock to each carrier is proportional to its measured decline in Suez use, scaled by a common factor chosen so that the decline in container capacity passing through Suez in the model's equilibrium matches the decline in the data. Conditional on these shocks, we examine what the model implies for outcomes that are not targeted: where the withdrawn capacity is redeployed, how U.S.-bound cargo is rerouted, and how freight charges and import quantities respond across markets that relied on Suez to different degrees before the crisis. Comparing these outcomes with the data provides a test of the model's mechanisms.

We construct the shocks in two steps. First, for each edge whose sea path passes through Suez or Bab el-Mandeb, we compute the shortest alternative route between the same ports on the ocean network of Section 3.3. Of the 200 affected service legs, 119 can be rerouted around the Cape of Good Hope, which increases their sailing requirements by 91 percent on average. The longer requirements imply the change in round-trip access $\widehat{\Phi}^{\text{full}}_{he} < 1$ that pool h would face on edge e if it withdrew from Suez entirely. The remaining 81 legs start or end within the Red Sea, cannot be rerouted between the same ports, and are left unchanged.

Second, the observed withdrawals determine how much of this full shock each carrier receives.

Maersk, MSC, and Hapag-Lloyd, which we track directly in GFW, cut their capacity passing Suez by 99, 90, and 100 percent between July–October 2023 and January–August 2024; for the remaining operators, we use the 63 percent fall in transits recorded by PortWatch. Letting d_{ht} denote the measured decline of pool h in month t, the shock is

$$\log \widehat{\Phi}_{he,t} = \min\{1, \delta^{RS} d_{ht}\} \log \widehat{\Phi}^{\text{full}}_{he},$$

where δ^{RS} is the common scale factor, estimated at 1.155. Rerouted edges also pay a higher exogenous cost reflecting the longer trip. All other parameters are held at their benchmark values, except that port congestion is shut down.

Results: Vessel reallocation. We now turn to the untargeted implications of the experiment. Because the aggregate decline in Suez use is matched by construction, the first question is where the vessel capacity withdrawn from Suez is redeployed. Figure 13 shows that both the data and the model shift capacity away from Suez and toward the Cape of Good Hope. The model therefore identifies the Cape as the principal alternative route observed in the data. Capacity through the Cape rises by 93 percent of the capacity that passed Suez before the event in the model, against 83 percent in the data.

Results: U.S. imports and freight charges. We next compare how U.S.-bound trade responded in the model and in the data, across import markets that relied on Suez to different degrees before the crisis. Panels (a), (b), and (e) of Figure 14 use Panjiva shipments linked to GFW vessel routes, with markets defined by origin country, last foreign port, and U.S. port of arrival. Panels (c), (d), and (f) use Census data, with markets defined by origin country, product, and customs district; because Census does not record vessels, we measure each market's reliance on Suez using route shares from Panjiva. For each outcome q and source $s \in \{D, M\}$ (data or model), we estimate

$$y^{s,q}_{nt} = \alpha^{s,q}_{n} + \Lambda^{s,q}_{nt} + \pi^{s}_{q} Z^{0}_{n} \,\mathrm{Post}_{t} + u^{s,q}_{nt}, \tag{5.1}$$

where $\alpha^{s,q}_{n}$ are market fixed effects, $\Lambda^{s,q}_{nt}$ time-varying controls, Z^{0}_{n} the market's pre-crisis Suez exposure, and Post_{t} equals one from January through August 2024, excluding November and December 2023 as adjustment months. We apply the same specification, markets, and weights to the data and to the model's baseline and post-disruption equilibria, so π^{s}_{q} measures the differential response of a fully exposed market in each. For each month, we solve for the equilibrium implied by that month's carrier withdrawals, so the model series in the figure trace the evolution of the shock rather than gradual adjustment over time.

We begin with cargo routing, an equilibrium outcome of the model. The model closely matches how cargo bound for the United States was rerouted across import markets (panels a and b): moving from zero to full pre-crisis exposure lowers the Suez route share by 98.9 percentage points in the data and 88.9 in the model, and raises the Cape share by 95.1 and 89.2. The model also captures the consequences

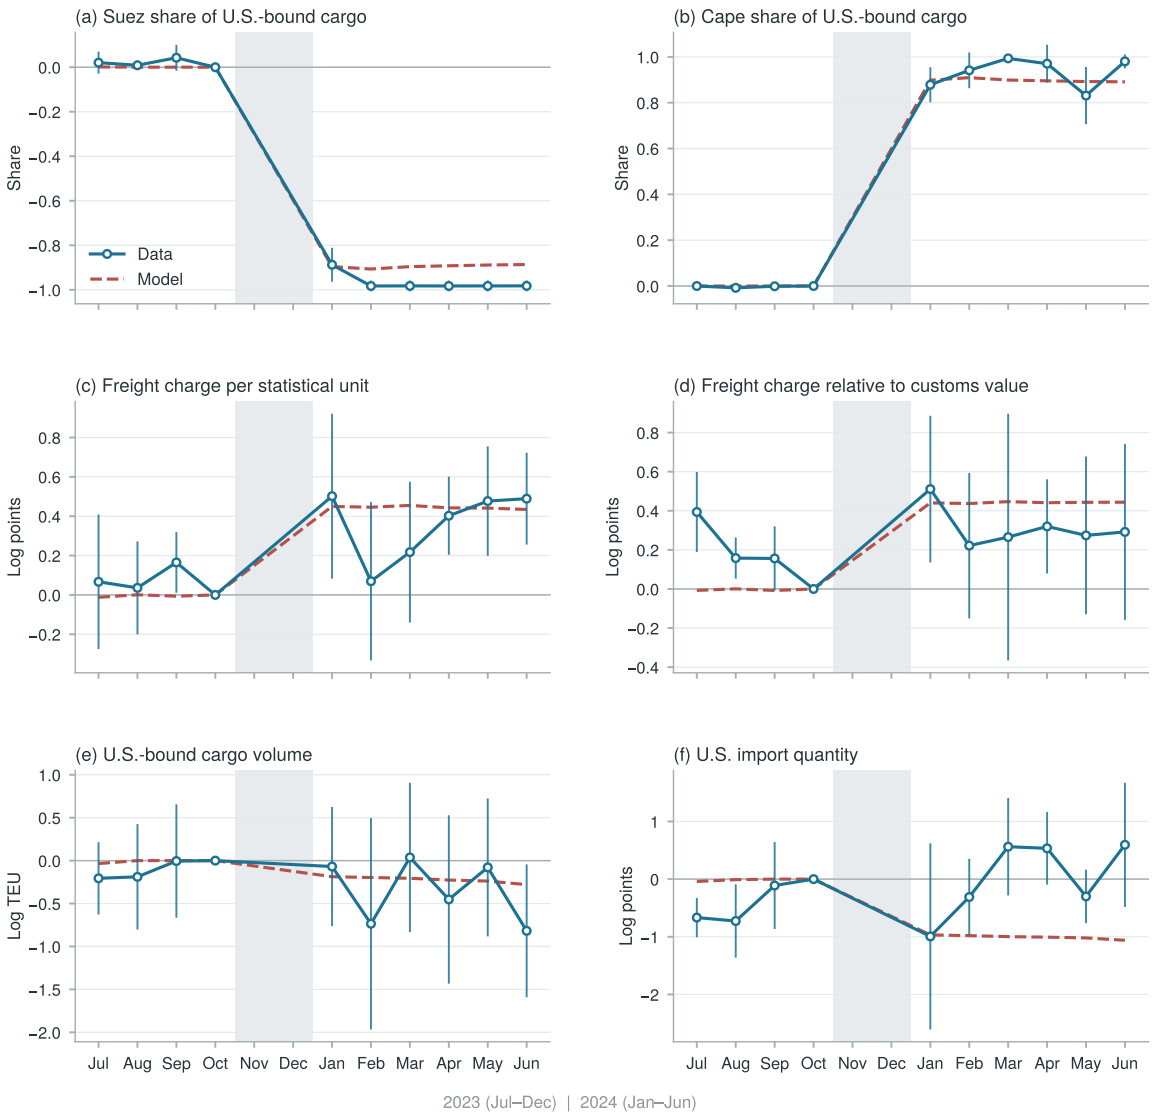


Figure 14. U.S. import responses to the Red Sea disruptions. Monthly coefficients on pre-crisis Suez exposure, relative to October 2023, for the data (blue, with 95 percent confidence intervals) and the model (red dashed). Panels (a), (b), and (e) use Panjiva shipments linked to GFW vessel routes; panels (c), (d), and (f) use U.S. Census data. Gray marks the November–December 2023 adjustment period, which is excluded. The figure shows months through June 2024.

of this rerouting. In the model, longer rotations tie up vessel capacity and raise freight rates, so freight charges should rise more in markets that relied more on Suez. The data show this pattern (panel c): moving from zero to full exposure raises log freight charges per unit by 0.223 in the data and 0.446 in the model, and the difference is not statistically significant ($p = 0.20$). Matched cargo volumes also decline similarly in exposed markets, by -0.297 (s.e. 0.267) log points in the data and -0.233 in the model (panel e). The model's predictions are larger than observed for the rise in freight charges common to all markets, 33.5 percent against 7.8, and for the responses of freight charges relative to customs value and of Census import quantities (panels d and f), although the latter two Census series exhibit differential pre-trends.

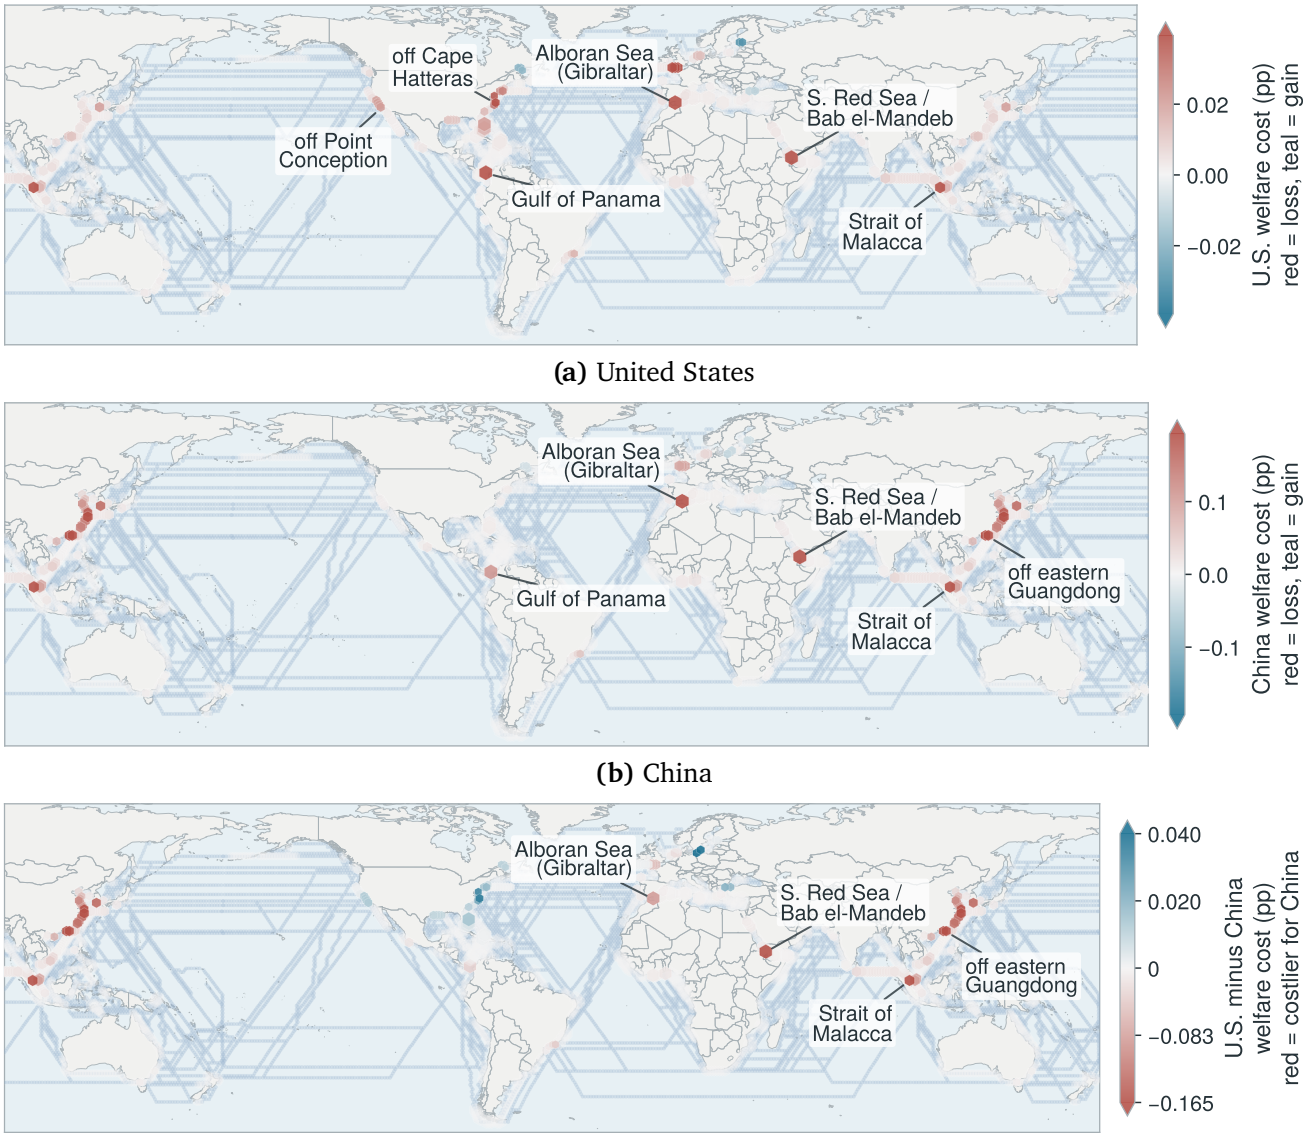
Overall, the episode supports the mechanism at the core of the model: when disruptions lengthen voyages, they tie up vessel capacity and raise shipping costs, more so for trade that relied on the disrupted route. A single event provides limited statistical power to discriminate among models, and the tests of Adão, Costinot and Donaldson (2025) reject, at the 5 percent level, neither the model ($p = 0.276$) nor a prediction of no change ($p = 0.073$). Still, the model captures the direction and cross-market pattern of the observed responses without targeting them. We next use it to study disruptions that have not occurred, where the same mechanism determines how geographic and capacity shocks propagate through shipping and trade.

6 Geographic Maritime Vulnerabilities

This section uses the model to quantify the welfare costs of geographic disruptions. As in the Red Sea disruptions studied in the previous section, the shocks depend on maritime geography: closing part of the ocean forces ships onto longer routes or leaves some routes without an alternative. We begin by closing one location at a time throughout the global shipping network, which maps where disruptions are most costly for each country. We then study how local vulnerabilities combine when several neighboring passages are disrupted at once. We focus on the waters around Taiwan and the South China Sea, which contain some of the most heavily used passages in the network, but also several nearby passages that traffic can switch to. Because these passages lie close to one another, a single disruption can affect several of them at once, making the region a natural setting to study combined disruptions and to ask whether heavy use of a passage translates into large welfare losses.

6.1 Global Geographic Disruptions

We begin by investigating the welfare consequences of geographic disruptions across the global shipping network. To do so, we subject every location along the network's sea routes to the same local disruption, one at a time, and compute the resulting welfare cost for each country. This exercise identifies the locations whose disruption is most costly, documents how their importance varies across countries, and



(a) United States

(b) China

(c) U.S.-minus-China welfare-cost difference

Figure 15. Welfare costs of local maritime disruptions. Panels (a) and (b) report the welfare cost of closing each cell for the United States and China; panel (c) reports the U.S. minus the Chinese cost, so teal marks closures that cost the United States more and red those that cost China more. Welfare costs are in log percentage points. Each panel has its own color scale, truncated at extreme values; in panel (c), positive and negative values are truncated separately.

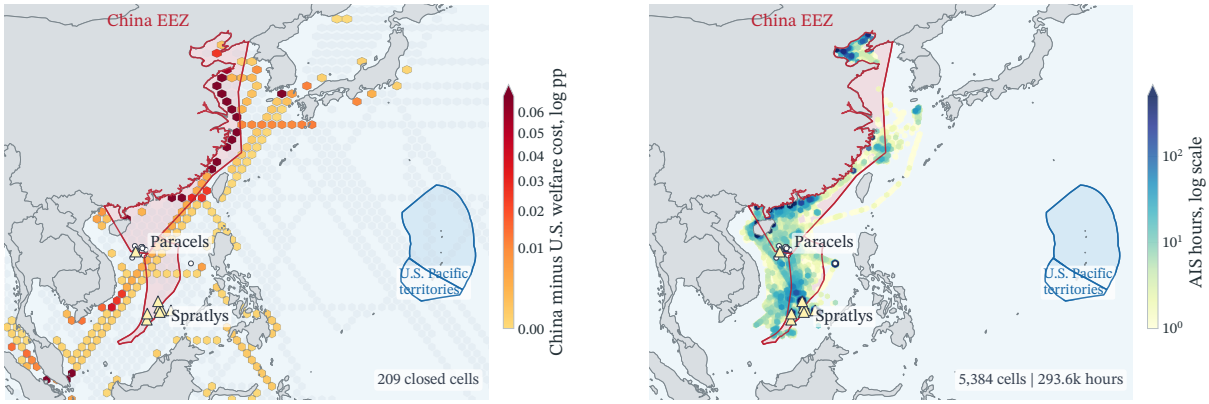
shows that this importance depends not only on how much trade passes through a location but also on how easily ships can avoid it.

The experiment works as follows. We divide the ocean into hexagonal cells about 100 kilometers across. We then close one cell at a time and reroute every affected port pair along the shortest remaining path. Rerouting lowers round-trip access $\widehat{\Phi}_{he}$ and raises the exogenous cost of the affected edges, by more where the detour is longer, and we solve for the new equilibrium. Because every closure removes an area of the same size, differences in welfare costs across cells reflect geography: how much trade passes through a cell and how easily ships can sail around it. In passages narrower than a cell, such as the southern Red Sea and the Strait of Malacca, closing a single cell closes the entire passage, and ships must reroute through other passages, around the Cape of Good Hope or through the Sunda Strait, for example. Lanes left without any route are removed from the network.

Results. Figure 15 maps the welfare cost of each local disruption for the United States and China. The two countries' vulnerabilities differ in both location and magnitude. U.S. losses are spread across distant chokepoints and the approaches to its own coasts: the largest, 0.133 log percentage points, results from a closure in the southern Red Sea near Bab el-Mandeb, followed by closures near Gibraltar and in the Strait of Malacca, and closures near Panama and off the U.S. East and West Coasts also impose sizable losses. China's losses are larger and more concentrated in East Asia: the largest, 0.611 log percentage points, results from a closure in the Strait of Malacca, and many of its costliest closures lie along the maritime approaches to Chinese ports.

Panel (c) reports the difference between the U.S. and Chinese welfare costs of each closure, which distinguishes how important a location is from how unevenly its disruption affects the two countries. Closures off the U.S. East Coast and in the Baltic are costlier for the United States than for China, while closures in the Red Sea and East Asia are costlier for China. Even at the chokepoints that matter most for both countries, such as the southern Red Sea, the approach to Gibraltar, and the Strait of Malacca, China loses more. The contrast is sharpest in the waters off China, shown in panel (a) of Figure 16: of the 209 cells closed there, 80 cost China more than the United States by over 0.001 log percentage points, and none cost the United States more by that margin. Outside Malacca, the largest gaps lie along China's coast, east of the Pearl River Delta and off Ningbo and Shanghai. Panel (b) shows that these waters overlap with the areas where vessels identified as part of China's maritime militia (Center for Advanced Defense Studies and Asia Maritime Transparency Initiative, 2021) were present in 2023. The overlap is consistent with these waters being of economic importance to China.

Some countries face much larger losses than either the United States or China. Across all closures, the largest welfare losses are 21.98 log percentage points for Yemen, 12.41 for Malaysia, 8.98 for Morocco, and 6.70 for French Polynesia, from closures in the southern Red Sea, the Strait of Malacca, near Gibraltar, and in the Gulf of Panama, respectively. These economies serve as transit hubs or ports of call, so their losses reflect cargo passing through their ports as well as their own trade. Taken together,



(a) China minus U.S. cost of closing each cell

(b) Militia vessels, 2023 AIS presence

Figure 16. Closures near China's coast and maritime-militia presence. Panel (a) reports China's minus the U.S. welfare cost of closing each cell, in log percentage points. Panel (b) reports the hours of 2023 AIS presence (Global Fishing Watch, 2025) of vessels on the maritime-militia list of Center for Advanced Defense Studies and Asia Maritime Transparency Initiative (2021), on a log scale. The red outline, labeled China EEZ, is drawn from Marine Regions boundary data (Flanders Marine Institute, 2023). Color scales are truncated at extreme values.

these results show that the welfare cost of disrupting a location is not simply a function of how much traffic passes through it, but also of whether ships can sail around it.

6.2 Disruptions at Multiple Chokepoints: Taiwan and the South China Sea

The previous subsection studies disruptions at one location at a time. In practice, disruptions can affect several nearby chokepoints at once. Ships then lose not only the closed routes but also the alternatives they would otherwise use, so the effect of closing one chokepoint depends on which others remain open. The waters around Taiwan and the South China Sea are a natural setting to study such disruptions, since several of the most heavily used chokepoints in the global shipping network lie close to one another. We first examine how closing several of them at once affects the availability and length of shipping routes. Because closures at several chokepoints affect countries' trade very unevenly, they can impose large losses on some countries at little cost to others. This asymmetry is central to the trade-off in Section 2.3, where a country considering a disruption weighs the loss it imposes on the target against the cost it bears itself. We measure the welfare side of both in a scenario drawn from public studies of a potential blockade of Taiwan (Lin et al., 2024*a*; Cancian, Cancian and Heginbotham, 2025), in which cargo around Taiwan is restricted while cargo to and from China continues to move. Finally, we show how restricting China-related cargo at other passages in the region raises the cost borne by China, and how it affects third countries.

Panel (a) of Figure 17 maps container shipping through eight straits in the region: the Taiwan Strait and the Luzon, Mindoro, Balabac, Malacca, Sunda, Makassar, and Lombok Straits. Traffic is highly concentrated: the Strait of Malacca alone carries 9.6 percent of all laden container traffic in the network,

Table 2. Route substitution under regional disruptions

Scenario	Affected service legs	Laden TEU (million)	Legs without an alternative	Mean distance increase (%)
2022 exercise-zone disruption	170	26.2	0	4.3
Taiwan passage unavailable	445	77.4	131	36.3
Broad regional zone	528	85.8	206	176.7

Notes: Affected legs are service legs whose baseline sea path enters the closed area, and laden TEU is their annual laden traffic. Legs without an alternative cannot connect the same ports after the closure. Mean distance increases are weighted by laden TEU across legs that can be rerouted.

and the Taiwan Strait 7.3 percent. How much each country relies on these straits, however, varies. For China, the Taiwan Strait carries far more traffic than Malacca: 27 percent of container traffic arriving at Chinese ports and 23 percent of traffic leaving them crosses the Taiwan Strait, compared with only 2 and 3 percent for Malacca.

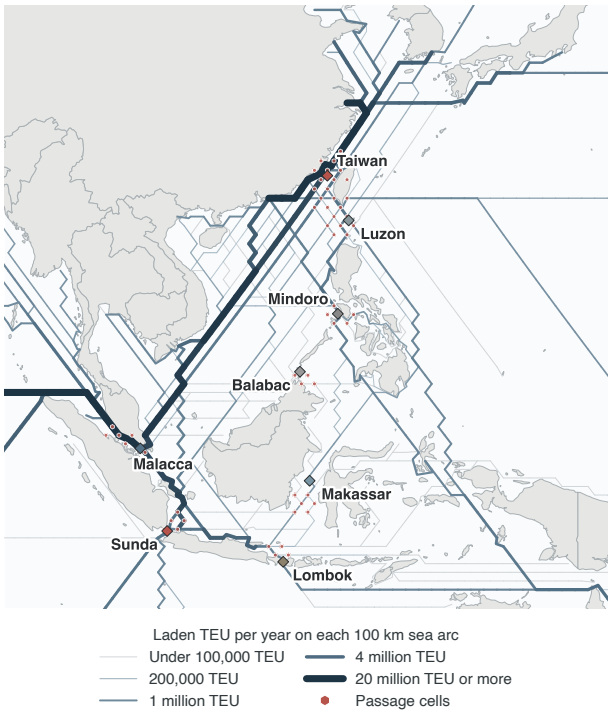
6.2.1 Regional Disruptions and Route Substitution

We consider three closures of increasing scope, each applying to all ships. The first closes the six maritime areas announced for military exercises around Taiwan in August 2022.[6] The second closes the Taiwan Strait while leaving neighboring straits open. The third, which we call the broad regional zone, extends the exercise areas by 50 kilometers and also closes the Luzon and Mindoro Straits and the eastern entrances to the central Philippine seas, following scenarios analyzed in public studies (Lin et al., 2024*a*; U.S. Department of Defense, 2024). For each closure, we hold the ports of every affected leg fixed and compute the shortest route that avoids the closed area. Legs without such a route lose their connection, since reconnecting their ports would require new shipping services. Panels (b) and (c) of Figure 17 illustrate this for the broad regional zone, showing the legs it affects, the detours of those that can be rerouted, and those that lose their route.

Table 2 reports the results. The 2022 exercise areas can be avoided at modest cost: all 170 affected legs, carrying 26.2 million laden TEU, can be rerouted, and their voyages lengthen by only 4.3 percent on average. Broader closures are far more disruptive. They cut off ports: closing the Taiwan Strait leaves 131 of 445 affected legs without an alternative, and the broad regional zone leaves 206 of 528 legs without one, mostly legs serving Taiwan's ports. They also force much longer detours on the legs that can still be rerouted, lengthening voyages by more than a third under the Taiwan Strait closure and nearly tripling them under the broad zone. How much traffic a closure affects thus measures its scale, but whether and how cheaply that traffic can be rerouted determines how disruptive it is.

[6] We construct the six areas from the coordinates in the August 2022 notices published by the PRC Ministry of National Defense/Xinhua and Taiwan's Maritime and Port Bureau.

(a) Baseline routes using the regional passages



(b) Disruption and affected legs

(c) Alternative routes and loss of route support

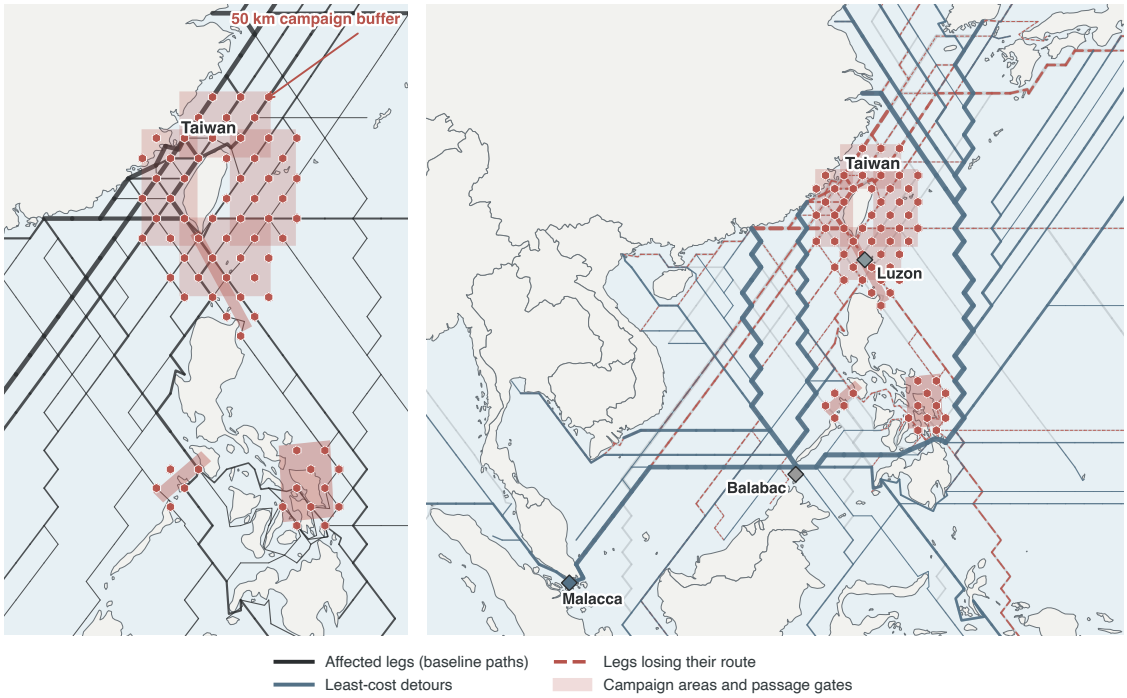


Figure 17. Regional exposure and route substitution. Laden container traffic on service legs. Panel (a) shows legs that cross the eight regional straits, marked in red. Panel (b) shows the broad regional zone and the legs it affects. Panel (c) distinguishes legs that can be rerouted, shown along their shortest detours, from legs that lose their route because one of their ports lies within the zone, shown dashed. Line width increases with traffic.

6.2.2 A Selective Disruption around Taiwan

The disruptions above restrict all ships alike, so they also impose costs on any country that might initiate them. Section 2.3 shows that, for the incentive to initiate a disruption, what matters is the loss imposed on the target relative to the cost borne by the initiating country. This ratio is larger when a disruption is selective, restricting some cargo while leaving other cargo unaffected. Closing the broad regional zone cuts off ports rather than only lengthening routes (Section 6.2.1), while nearby straits allow unrestricted cargo to keep moving. We therefore study a selective version of this scenario, drawing on public studies of a potential blockade of Taiwan (Lin et al., 2024*a*; Cancian, Cancian and Heginbotham, 2025): Taiwan's international lanes and the surrounding zone are closed to cargo that neither originates in nor is bound for China, while cargo to and from China continues to move. We quantify the welfare costs this imposes on each country, and Section 6.2.3 then examines how restrictions on China-related cargo at other straits would change the cost borne by China.

To implement this disruption, we split cargo into two types: China-related cargo, which originates in or is bound for China, and all other cargo.[7] The two types travel on the same ships and pay the same freight rates, but face different restrictions. China-related cargo moves freely. Other cargo cannot use any lane to or from Taiwan's ports, and must sail around the zone; if no route around it exists, that cargo cannot be shipped. A service that is no longer allowed to carry any cargo closes, and its vessels are redeployed. We measure the share of each type of cargo on every lane from the traffic assignment of Section 3.1.

Table 3 reports the welfare costs. Losses are concentrated in nearby economies. Taiwan loses 10.39 log percentage points, as its trade with partners other than China has no permitted route, followed by Hong Kong (7.18), Vietnam (3.27), Singapore (2.07), South Korea (1.76), and the Philippines (1.47). The United States loses 0.015. China's welfare rises slightly, by 0.36 log percentage points, as its cargo gains access to the routes and vessel capacity released by restricted trade. In terms of the trade-off in Section 2.3, the disruption imposes a large loss on the target, while its welfare cost to the initiating country through the shipping channel is negative.

Figure 18 shows how restricted cargo adjusts. Of the restricted cargo on affected legs, 41 percent loses its route entirely. The remainder detours around the zone, but the longer voyages raise shipping costs enough that only 27 percent of it is still shipped in equilibrium. These calculations isolate the international container-shipping channel. In this respect, our analysis is related to Coşar and Thomas (2021), who also map disruptions to Southeast Asian waterways into changes in sailing distances and trade. A full assessment of a blockade or military conflict would also account for channels outside the model, such as tanker and bulk shipping, semiconductors, inventories, energy, finance, and military escalation (Verschuur, Lumma and Hall, 2025; Hart et al., 2026; Vest, Kratz and Goujon, 2022; Lin et al., 2024*b*,*a*; Cancian, Cancian and Heginbotham, 2025).

[7] Hong Kong is treated as a separate economy, as it is in the trade data, so its trade with countries other than mainland China is restricted.

Table 3. Welfare costs of a selective disruption around Taiwan

		With restrictions on China-related cargo in	
Economy	Disruption	the Taiwan Strait	all eight straits
Taiwan	10.393	10.445	10.413
Hong Kong	7.179	6.611	5.164
Vietnam	3.274	2.441	0.816
Singapore	2.070	1.535	−0.690
South Korea	1.764	1.470	1.198
Philippines	1.468	1.395	1.008
Japan	0.565	0.479	0.343
China	−0.363	−0.234	0.146
United States	0.015	−0.019	−0.083

Notes: Welfare costs in log percentage points relative to the baseline; positive values are losses. The selective disruption restricts cargo with neither origin nor destination in China on Taiwan's international lanes and in the broad regional zone. The last two columns additionally restrict China-related cargo in the Taiwan Strait, or in all eight straits.

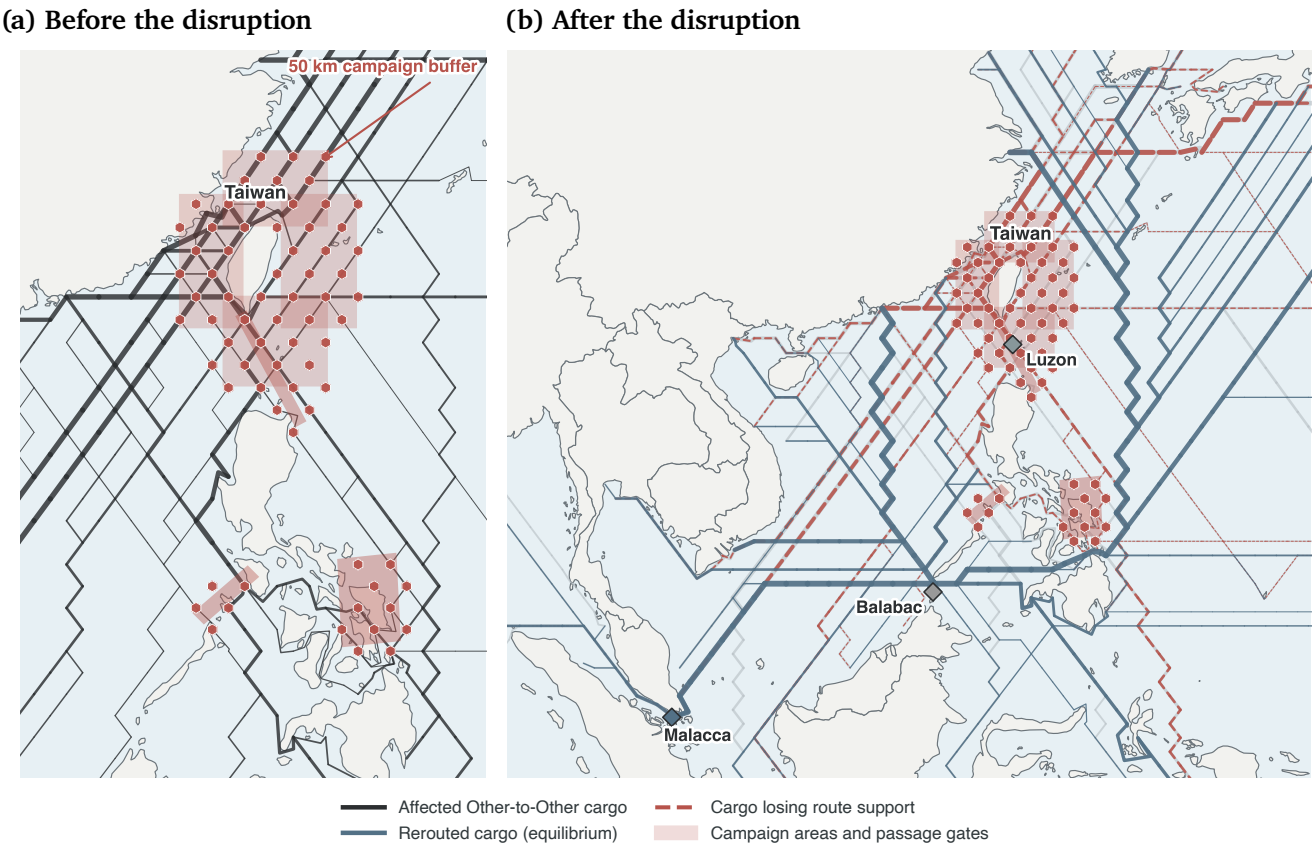


Figure 18. Cargo diversion under the selective disruption. The figure shows trade that neither originates in nor is bound for China, the trade restricted by the disruption. Panel (a) shows this trade on its baseline routes through the affected area. Panel (b) shows how it is diverted: solid lines show trade still shipped along detours around the restricted area, and dashed lines show trade that can no longer be shipped, including Taiwan's trade with countries other than China. Line width increases with the value of trade.

6.2.3 Restrictions on China-Related Cargo at Other Straits

Proposition 2.2 implies that the lower the cost a disruption imposes on the initiating country relative to the loss it imposes on the target, the lower the threshold above which it becomes attractive. This cost is not fixed by the disruption alone: it also depends on how other countries respond. One possible response is to restrict the initiating country's own cargo elsewhere in the network. How effective such restrictions are depends on maritime geography, since cargo restricted at one chokepoint can shift to others nearby. The waters around Taiwan and the South China Sea, with several busy straits close to one another, are a natural setting to study this question. Section 6.2.2 shows that the selective disruption imposes large losses on Taiwan and nearby economies while, through the shipping channel, imposing no welfare cost on China. Starting from that equilibrium, we restrict China-related cargo at combinations of the eight straits in the region and compute how much each combination raises China's welfare cost. Because the effect of each restriction depends on which other straits are also restricted, we evaluate all combinations jointly. We also examine how these restrictions affect third countries.

Each restriction applies to China-related cargo at a given strait, while all other cargo is unaffected. Restricted cargo must sail around the strait where a route exists, and cannot be shipped where none does. As before, longer detours raise shipping costs and tie up vessel capacity. Ports located on the straits require one further choice. For Malacca, Sunda, and the other southern straits, the restriction also covers China-related cargo to and from ports on the strait, such as Port Klang and Jakarta. For the Taiwan Strait, it covers only cargo sailing through the strait, so that China's own ports along it, such as Xiamen and Fuzhou, remain open. With eight straits, there are 255 possible combinations of restrictions. We evaluate all of them and, for each number of restricted straits, report the combination that most raises China's welfare cost relative to the selective disruption alone.

Table 4 reports the results. China ends up worse off than without the selective disruption only when several straits are restricted at once. Restricting Malacca alone, the single strait with the largest effect, raises China's cost by 0.16 log percentage points, and restricting the Taiwan Strait by 0.13, both less than China's gain from the selective disruption. Figure 19 shows why. Restricting the Taiwan Strait cuts no route, as China-related cargo detours east of Taiwan and 80 percent of it is still shipped. Restricting the southern straits cuts off the ports along them, which handle 45 percent of the affected cargo, but the rest detours south of Java, and 87 percent of it is still shipped. Only 32 of the 255 combinations leave China worse off than without the disruption, and all of them include the Taiwan Strait, Malacca, and Sunda, the main alternative to Malacca. Additional straits add little: restricting all eight leaves China with a net welfare loss of 0.15 log percentage points.

These restrictions also affect third countries, often more than China. Restricting all eight straits raises China's welfare cost by 0.51 log percentage points but Yemen's by 2.05, reflecting cargo that passes through its ports. Other economies gain relative to the selective disruption alone, as the restrictions partly offset its effects. Throughout, we take these restrictions as given and do not assess their enforcement or how their prospect would affect the initial decision.

Table 4. Restrictions on China-related cargo at other straits

No.	Straits restricted	Added welfare cost			China's total cost	Largest added cost, third country
		China	Taiwan	U.S.		
1	Malacca	0.162	−0.112	−0.040	−0.201	Tanzania (1.45)
2	Taiwan Strait, Malacca	0.298	−0.023	−0.071	−0.065	Tanzania (1.50)
3	Taiwan Strait, Malacca, Sunda	0.404	−0.025	−0.084	0.041	Yemen (1.71)
4	Taiwan Strait, Malacca, Sunda, Lombok	0.459	−0.001	−0.090	0.097	Yemen (1.98)
5	Taiwan Strait, Luzon, Malacca, Sunda, Lombok	0.500	0.029	−0.098	0.138	Yemen (2.06)
6	Taiwan Strait, Luzon, Mindoro, Malacca, Sunda, Lombok	0.508	0.022	−0.098	0.146	Yemen (2.06)
7	Taiwan Strait, Luzon, Mindoro, Malacca, Sunda, Makassar, Lombok	0.509	0.020	−0.098	0.146	Yemen (2.05)
8	All eight straits	0.509	0.020	−0.098	0.146	Yemen (2.05)

Notes: Welfare costs in log percentage points; positive values are losses. For each number of restricted straits, the table reports the combination that most raises China's welfare cost, with the selective disruption in place. Added costs are relative to the selective disruption alone; China's total cost is relative to the baseline. Restricted China-related cargo detours around a strait where a route exists and cannot be shipped otherwise.

Together, these results show how maritime geography shapes vulnerability to disruptions. Heavily used chokepoints need not be the most costly to disrupt when nearby alternatives exist, and closing several neighboring chokepoints at once matters because it removes those alternatives. Geography also allows disruptions to be selective. In the trade-off of Section 2.3, a selective disruption can impose large losses on some countries at little cost to the initiating country through the shipping channel, while restrictions on the initiating country's cargo make it worse off only when they cover several straits at once, and they impose substantial costs on third countries.

7 Maritime Supply Vulnerability

Global shipping requires not only reliable routes but also the vessels to sail them. Much of this capacity is built and operated abroad, so a country's trade can also be exposed to disruptions in the availability of foreign vessels, whether because operators withdraw them or because policy restricts their use. Two features shape the consequences of such disruptions. First, construction origin and operator control are distinct, so restrictions based on where vessels were built and restrictions based on who operates them can affect different parts of the fleet. Second, unlike the geographic disruptions of Section 6, routes may remain open; what would change is the supply of vessels, and the remaining fleet may need to be redeployed to fill the gap. These features raise two questions: how costly may such disruptions be, and

(a) Taiwan Strait **(b) Malacca, Sunda, Makassar, and Lombok**

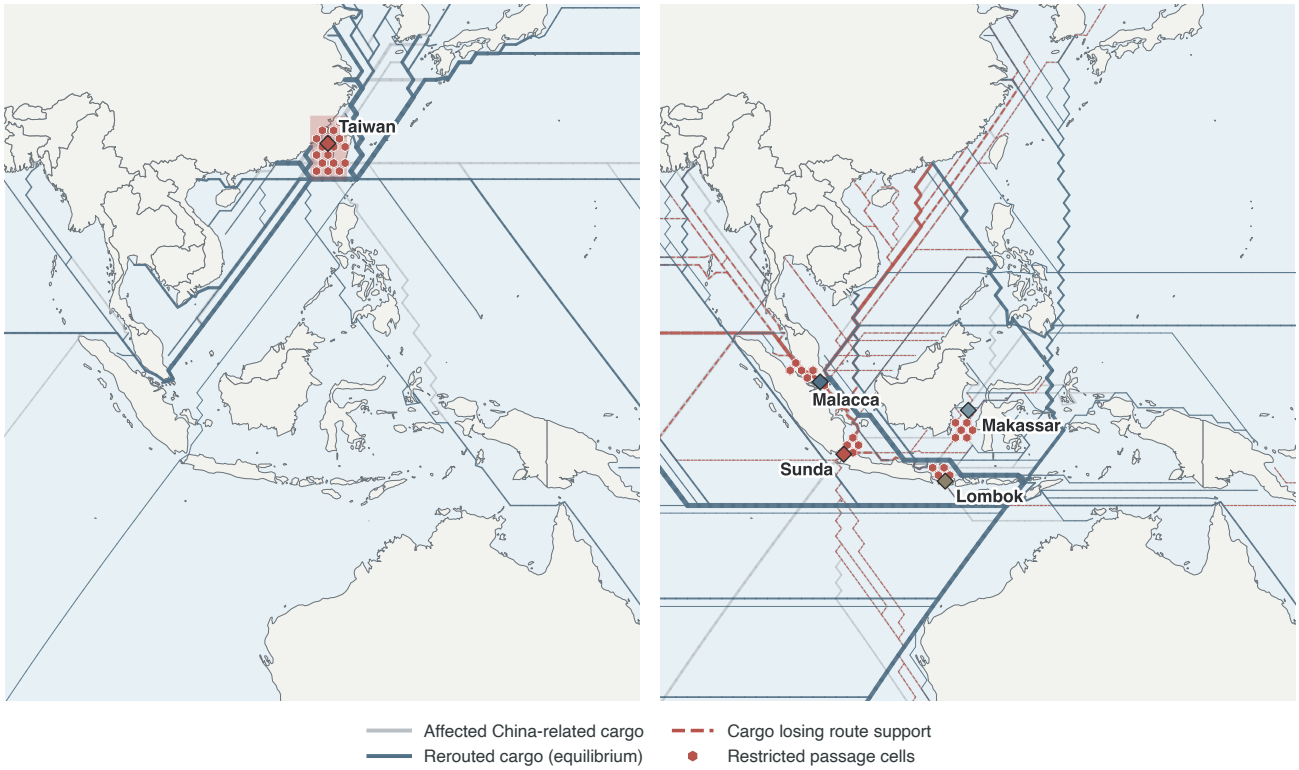


Figure 19. Cargo diversion under restrictions on China-related cargo. Panel (a) restricts the Taiwan Strait; panel (b) restricts Malacca, Sunda, Makassar, and Lombok. Gray lines show China-related cargo on affected legs along their baseline routes. Blue lines show the cargo still shipped along detours in equilibrium, and dashed red lines show cargo on legs that lose their route. Hexagons mark the restricted straits. Line width increases with the value of cargo.

to what extent may policies that reduce dependence on foreign capacity lower that cost.

We study these questions in the context of U.S. dependence on Chinese shipping capacity, the focus of recent U.S. policy, including the investigation of China's maritime and shipbuilding sectors and the fees announced in April 2025 on Chinese vessel operators and operators of Chinese-built ships (USTR, 2025). The distinction between construction and operation is particularly relevant here: most Chinese-built capacity is operated by carriers from other countries (Section 3.4). We first quantify the welfare effects of reducing the availability of Chinese-operated or Chinese-built capacity on routes serving the United States, both in the United States and abroad. We then use the model to compare the cost of the announced fees in normal times with the losses they could avoid if this capacity were disrupted.

7.1 Disruptions to Shipping Capacity

We study disruptions to shipping capacity through three counterfactual experiments on the international shipping lanes with a U.S. origin or destination. In each, the affected capacity becomes largely unavailable on these lanes: we lower its round-trip access Φ_{he} to 10 percent of its baseline value, while the physical network and total fleet remain unchanged. The experiments differ only in which vessels are affected. The first restricts capacity operated by Chinese carriers, wherever their vessels were built. The second restricts Chinese-built vessels, whoever operates them. The third restricts capacity that is either Chinese-operated or Chinese-built. In each case, vessel deployment, freight rates, and trade adjust in equilibrium.

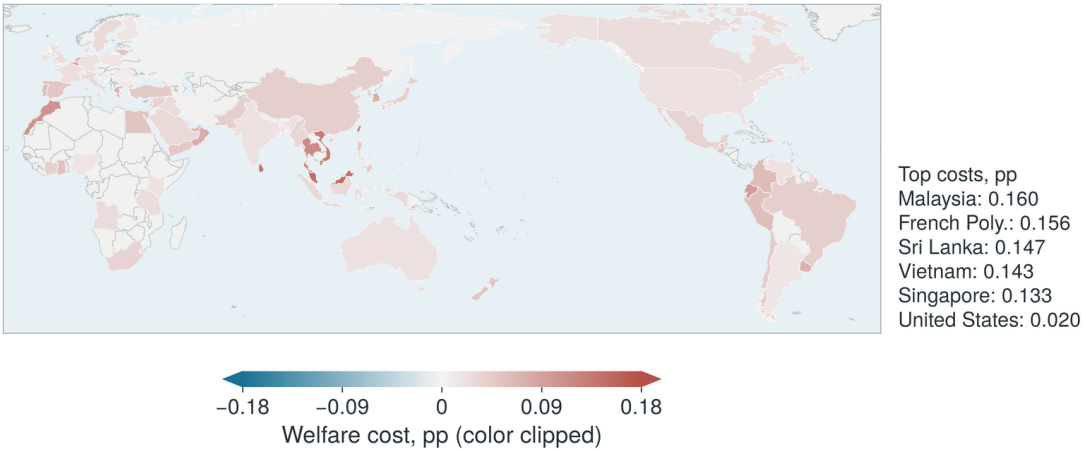
Figure 20 reports the welfare effects. Each panel maps the welfare cost for every country in one experiment: panel (a) restricts Chinese-operated capacity, panel (b) Chinese-built capacity, and panel (c) both. Although the first two restrictions affect different vessels, their costs for the United States are similar, 0.020 and 0.021 log percentage points, while restricting both costs 0.035. In all three experiments, the remaining fleet is redeployed to serve U.S. routes, which limits the losses but raises freight rates elsewhere in the network.

The losses are not concentrated in the United States. Because shipping routes compete for the same fleet, redeploying vessels toward U.S. routes raises shipping costs for other countries. Restricting both Chinese-operated and Chinese-built capacity, for example, imposes losses several times larger than the U.S. loss on economies in Southeast and South Asia, such as Malaysia, Sri Lanka, Vietnam, and Singapore, with smaller losses extending to Latin America, Africa, and Oceania; no economy gains.[8] The same economies bear the largest losses in all three experiments.

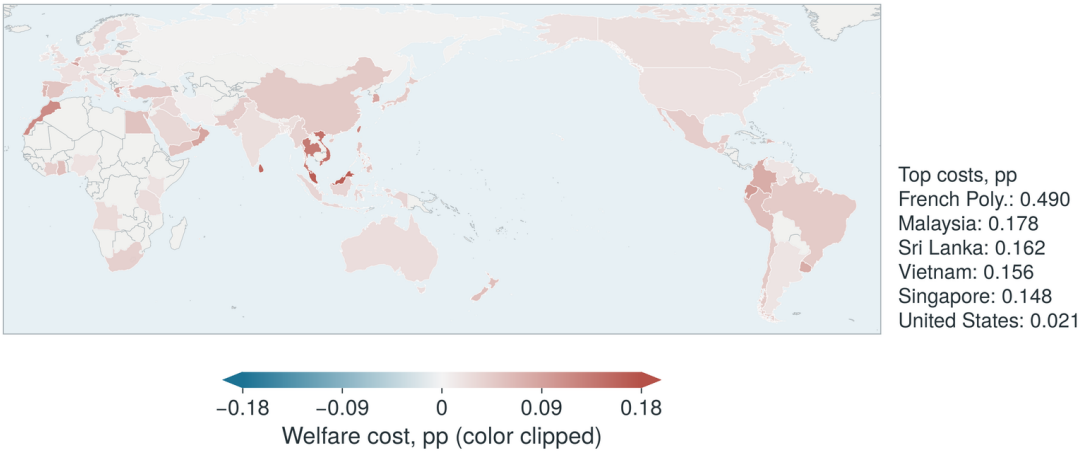
7.2 Prospective Policies to Reduce Maritime Dependence

Section 7.1 shows that the United States would lose if Chinese-operated or Chinese-built capacity became largely unavailable on its routes. A natural question is whether policy can limit such losses by lowering

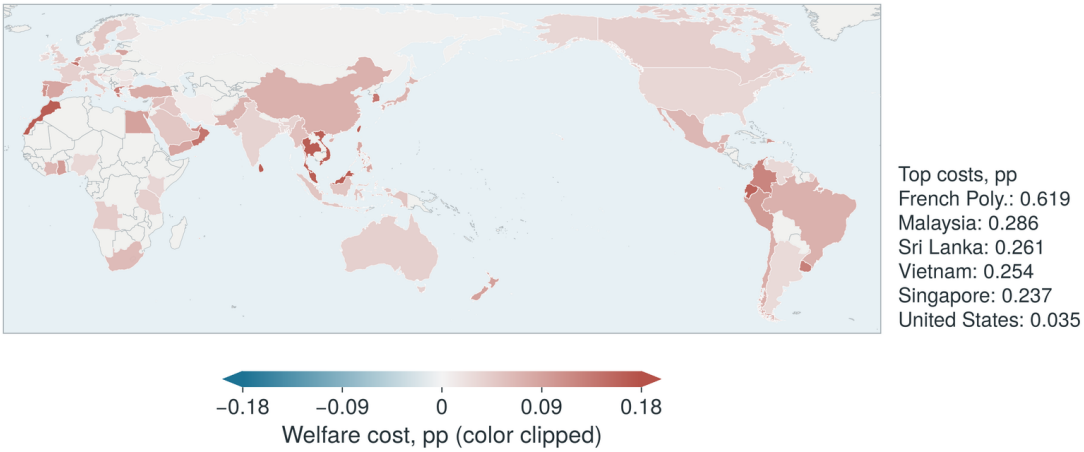
[8]The largest loss, for French Polynesia, reflects cargo on services between the U.S. East Coast and Oceania that call at Papeete.



(a) Chinese-operated vessels



(b) Chinese-built vessels



(c) Chinese-operated or Chinese-built vessels

Figure 20. Welfare effects of disruptions to shipping capacity. Panels (a), (b), and (c) restrict Chinese-operated capacity, Chinese-built capacity, and both, respectively, on international lanes serving the United States. Red indicates welfare losses and teal gains, in log percentage points. The three maps share one color scale, truncated at extreme values; the lists beside each map report the five largest losses and the U.S. loss.

Table 5. Shipping fees and U.S. welfare

Affected vessels	Policy	Normal times	Disruption	Additional loss
Chinese-operated	No fees	0.0000	0.0203	0.0203
	Fees	0.0040	0.0219	0.0179
Chinese-built	No fees	0.0000	0.0208	0.0208
	Fees	0.0040	0.0240	0.0200
Chinese-operated or Chinese-built	No fees	0.0000	0.0352	0.0352
	Fees	0.0040	0.0367	0.0327

Notes: U.S. welfare costs in log percentage points relative to the baseline without fees or disruption; positive values are losses. Fees are the rates previously scheduled for April 2026. Each disruption lowers the round-trip access of the affected vessels to 10 percent of its baseline value on international lanes serving the United States. The additional loss is the cost with the disruption minus the cost in normal times, under the same policy.

reliance on this capacity in advance. A policy of this kind works like insurance. It may raise shipping costs in normal times, by moving cargo away from the vessels that would otherwise carry it, but it may also lower losses if a disruption occurs, by reducing the trade that depends on the capacity that is lost. Whether it pays off depends on the size of these two effects and on how likely a disruption is. The fees announced by the United States in April 2025 (USTR, 2025) offer a concrete setting for measuring them. They apply to the same two groups of vessels studied above, those operated by Chinese entities and those built in China, and they are set at specific rates on vessels entering U.S. ports. Their design thus maps directly into the disruptions of Section 7.1, so we can quantify this trade-off for the fees as announced.

The fees were suspended in November 2025 and remain suspended. We study them at the rates that had been scheduled for April 2026. Vessels owned or operated by Chinese entities would pay $80 per net ton each time they enter the United States, and other Chinese-built vessels the higher of $23 per net ton or $153 per container discharged. In the model, the fees act as a tax on the freight these vessels earn on U.S. routes. For every vessel entering the United States, we compute its fee as a share of the freight it earns on the leg into the country.[9] We then aggregate these shares by operator country and lane, so the fee each country's carriers face on a lane reflects how many of their vessels there are Chinese-operated or Chinese-built. Shippers pay the market freight rate q_e, while carriers keep q_e minus their fee. Carriers therefore shift the affected capacity away from U.S. routes, and freight rates, vessel deployment, and trade adjust in equilibrium.

To measure both sides of the trade-off, we compare the economy with and without the fees, in normal times and when a disruption occurs. We do so for each of the three disruptions of Section 7.1, applying the same disruption with and without the fees. Table 5 reports U.S. welfare in these four cases, relative to the baseline without fees or disruption. The first column shows the cost of the fees in normal times, and the second the total cost when the disruption occurs. The third column shows the difference between the two, the additional loss caused by the disruption under each policy. Fees that work as insurance would raise the cost in the first column but lower the additional loss in the third.

[9]This requires imputing vessels' net tonnage, containers discharged, and freight receipts, the latter anchored to U.S. Census freight charges.

The fees do work as insurance. They cost the United States 0.004 log percentage points in normal times, and they lower the additional loss from each disruption. Under the operator disruption, for example, the additional loss falls from 0.020 to 0.018. Two forces lie behind this. The fees have already moved some trade away from the affected vessels, so losing those vessels matters less. And once these vessels are disrupted, they carry little cargo, so the fees themselves cost less.

In all three cases, however, the reduction in the additional loss is smaller than the cost of the fees in normal times. U.S. welfare is therefore lower with the fees whether or not a disruption occurs. Under the operator disruption, the United States loses 0.022 log percentage points with the fees, compared with 0.020 without them. In the model, the fees therefore lower expected U.S. welfare at any probability of disruption. The calculation holds the fleet fixed and does not capture effects of the fees on the likelihood of a disruption or on shipbuilding.

8 Conclusion

Dependence on foreign-controlled inputs is a central concern in debates about economic security. Maritime transportation makes this concern particularly relevant: countries rely on routes beyond their borders and on vessels built and operated abroad to conduct international trade. This paper studies when that dependence becomes a source of economic vulnerability. In doing so, we distinguish the welfare consequences of losing access to routes and vessel capacity from the economic incentives to impose those restrictions.

We develop an analytical framework linking trade exposure, shipping-cost responses, and equilibrium price adjustment to welfare, and combine global shipping and trade data with a spatial model featuring endogenous routing and vessel deployment. Our findings show that exposure need not imply vulnerability. Lanes carrying more trade are not necessarily harder to bypass, and the cost of disrupting a location depends on whether ships can sail around it. U.S. losses are spread across distant chokepoints, while China's are larger and concentrated in East Asia, and the largest losses fall on small economies near major chokepoints. Disruptions that affect several nearby chokepoints at once, as around Taiwan and the South China Sea, remove these alternatives: a disruption there that restricts only some countries' cargo can impose large losses on them at little cost to others, while restrictions on the initiator's cargo make it worse off only when they cover several straits at once, and they also affect third countries. Dependence on vessels matters as well. Because routes compete for the same fleet, a withdrawal of Chinese-operated or Chinese-built vessels from routes serving the United States imposes larger losses on several Asian economies than on the United States itself. These outcomes reflect how maritime geography and vessel capacity interact: longer voyages require additional capacity, and its reallocation changes shipping costs throughout the network.

The findings have implications for policies intended to reduce foreign dependence. Greater reliance on alternative capacity may limit disruption losses but make shipping more expensive even when no

disruption occurs. For the fees announced in 2025, the model implies that the reduction in disruption losses is smaller than the cost in normal times. Evaluating such policies therefore requires comparing their costs and benefits across both states of the world and accounting for their effects on other countries. Our analysis focuses on adjustment through existing routes and vessel capacity. Extending the analysis to vessel investment, shipbuilding, and new shipping services would allow us to study how policies reshape maritime dependence over time and alter the incentives to disrupt shipping.

References

Adamopoulos, Tasso, and Fernando Leibovici. 2025. "Trade Risk and Food Security." Federal Reserve Bank of St. Louis Working Paper 2024-004.

Adão, Rodrigo, Arnaud Costinot, and Dave Donaldson. 2025. "Putting Quantitative Models to the Test: An Application to the U.S.–China Trade War." *Quarterly Journal of Economics*, 140(2): 1471–1524.

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

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

Allen, Treb, Simon Fuchs, and Woan Foong Wong. 2025. "Evaluating Transportation Improvements Quantitatively: A Primer." Federal Reserve Bank of Atlanta Working Paper 2025-13.

Antràs, Pol. 2020. "Conceptual Aspects of Global Value Chains." *The World Bank Economic Review*, 34(3): 551–574.

Antràs, Pol, and Davin Chor. 2022. "Global Value Chains." In *Handbook of International Economics*. Vol. 5, 297–376. Elsevier.

A.P. Moller–Maersk. 2023. "Maersk Operations through the Red Sea / Gulf of Aden." *Customer advisory*.

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

Arslanalp, Serkan, Seung Mo Choi, Parisa Kamali, Robin Koepke, Matthew McKetty, Michele Ruta, Mario Saraiva, Alessandra Sozzi, and Jasper Verschuur. 2025. "Nowcasting Global Trade from Space." International Monetary Fund IMF Working Paper 2025/093.

Bachmann, Rüdiger, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick. 2024. "What If? The Macroeconomic and Distributional Effects for Germany of a Stop of Energy Imports from Russia." *Economica*, 91(364): 1157–1200.

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

Baldwin, David A. 1985. *Economic Statecraft.* Princeton, NJ:Princeton University Press.

Baldwin, Richard, Rebecca Freeman, and Angelos Theodorakopoulos. 2023. "Hidden Exposure: Measuring US Supply Chain Reliance." *Brookings Papers on Economic Activity*, 54(2): 79–167.

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

Barbieri, Katherine. 1996. "Economic Interdependence: A Path to Peace or a Source of Interstate Conflict?" *Journal of Peace Research*, 33(1): 29–49.

Barwick, Panle Jia, Myrto Kalouptsidi, and Nahim Bin Zahur. 2024. "Industrial policy: lessons from shipbuilding." *Journal of Economic Perspectives*, 38(4): 55–80.

Barwick, Panle Jia, Myrto Kalouptsidi, and Nahim Bin Zahur. 2025. "Industrial policy implementation: Empirical evidence from China's shipbuilding industry." *Review of Economic Studies*, 92(6): 3611–3648.

Becko, John Sturm. 2024. "A Theory of Economic Sanctions as Terms-of-Trade Manipulation." *Journal of International Economics*, 150: 103898.

Behrens, Kristian, and Pierre M. Picard. 2011. "Transportation, Freight Rates, and Economic Geography." *Journal of International Economics*, 85(2): 280–291.

Bezerra, P., and Alex Braithwaite. 2019. "Codebook for the Militarized Interstate Dispute Location (MIDLOC-A/I) Dataset, v2.1." *Correlates of War Project*, `https://correlatesofwar.org/data-sets/midloc/`.

Bianchi, Javier, and César Sosa-Padilla. 2024. "On Wars, Sanctions, and Sovereign Default." *Journal of Monetary Economics*, 141: 62–70.

Blackwill, Robert D., and Jennifer M. Harris. 2016. *War by Other Means: Geoeconomics and Statecraft.* Cambridge, MA:Harvard University Press.

Bonadio, Barthélémy. 2024. "Ports vs. Roads: Infrastructure, Market Access and Regional Outcomes." CESifo CESifo Working Paper 11383.

Bonadio, Barthélémy, Zhen Huo, Andrei A. Levchenko, and Nitya Pandalai-Nayar. 2021. "Global Supply Chains in the Pandemic." *Journal of International Economics*, 133: 103534.

Braithwaite, Alex. 2010. "MIDLOC: Introducing the Militarized Interstate Dispute Location Dataset." *Journal of Peace Research*, 47(1): 91–98.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. "Geography, transportation, and endogenous trade costs." *Econometrica*, 88(2): 657–691.

Brockhaus, Carsten Philipp, Julian Hinz, and Charles Serfaty. 2026. "Navigating Shocks: The Ripple Effects of Shipping Route Closures." Banque de France Working Paper Series 1057.

Brooks, Leah, Nicolas Gendron-Carrier, and Gisela Rua. 2018. "The local impact of containerization." *Finance and Economics Discussion Series*, 45.

Cancian, Mark F., Matthew F. Cancian, and Eric Heginbotham. 2025. "Lights Out? Wargaming a Chinese Blockade of Taiwan." Center for Strategic and International Studies.

Census Bureau. 2023–2024. "USA Trade Online." *https://usatrade.census.gov*, Last accessed on August 2024.

Center for Advanced Defense Studies, and Asia Maritime Transparency Initiative. 2021. "Pulling Back the Curtain on China's Maritime Militia." C4ADS and Center for Strategic and International Studies.

Chowdhry, Sonali, Inga Heiland, and Hendrik Mahlkow. 2026. "Quantitative Trade with Ships." Kiel Institute for the World Economy Kiel Working Paper 2314.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger. 2024. "A Theory of Economic Coercion and Fragmentation." National Bureau of Economic Research Working Paper 33309.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger. 2026. "A Framework for Geoeconomics." *Econometrica*, 94(1): 105–136.

CMA CGM. 2023. "CMA CGM Group's Preventive Measures for Navigation in the Red Sea." *Customer advisory*.

Coşar, A Kerem, and Banu Demir. 2018. "Shipping inside the box: Containerization and trade." *Journal of International Economics*, 114: 331–345.

Coşar, A. Kerem, and Benjamin D. Thomas. 2021. "The geopolitics of international trade in Southeast Asia." *Review of World Economics*, 157(1): 207–219.

Costinot, Arnaud, and Andrés Rodríguez-Clare. 2014. "Trade Theory with Numbers: Quantifying the Consequences of Globalization." In *Handbook of International Economics*. Vol. 4, 197–261. Elsevier.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum. 2008. "Global Rebalancing with Gravity: Measuring the Burden of Adjustment." *IMF Staff Papers*, 55(3): 511–540.

Devetak, Mitja, Jasper Verschuur, and Peter Klimek. 2026. "Adaptive Rerouting Reshapes Impacts of Maritime Chokepoint Disruptions." *arXiv preprint 2606.13431*.

Do, Anh D., Sharat Ganapati, Woan Foong Wong, and Oren Ziv. 2025. "Transshipment Hubs, Trade, and Supply Chains." Centre for Economic Policy Research CEPR Discussion Paper 20731.

Dunn, Jason, and Fernando Leibovici. 2024. "Navigating the Waves of Global Shipping: Drivers and Aggregate Implications." Federal Reserve Bank of St. Louis.

Eaton, Jonathan, and Maxim Engers. 1992. "Sanctions." *Journal of Political Economy*, 100(5): 899–928.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, 70(5): 1741–1779.

Elliott, Matthew, Benjamin Golub, and Matthew V. Leduc. 2022. "Supply Network Formation and Fragility." *American Economic Review*, 112(8): 2701–2747.

Fajgelbaum, Pablo D., Amit Khandelwal, Theodore Naff, and Edouard Schaal. 2026. "Strategic Trade Infrastructure: Understanding China's Belt and Road Initiative." Manuscript in preparation; presented at the NBER Summer Institute.

Fajgelbaum, Pablo D, and Edouard Schaal. 2020. "Optimal transport networks in spatial equilibrium." *Econometrica*, 88(4): 1411–1452.

Fan, Jingting, and Wenlan Luo. 2020. "A Tractable Model of Transshipment." Working Paper.

Farrell, Henry, and Abraham L. Newman. 2019. "Weaponized Interdependence: How Global Economic Networks Shape State Coercion." *International Security*, 44(1): 42–79.

Farrell, Henry, and Abraham L. Newman. 2023. *Underground Empire: How America Weaponized the World Economy.* New York:Henry Holt and Company.

Feyrer, James. 2021. "Distance, Trade, and Income – The 1967 to 1975 Closing of the Suez Canal as a Natural Experiment." *Journal of Development Economics*, 153: 102708.

Flanders Marine Institute. 2023. "Maritime Boundaries Geodatabase: Maritime Boundaries and Exclusive Economic Zones (200NM), Version 12." *Marine Regions*, available online at https://www.marineregions.org/. https://doi.org/10.14284/632.

Fuchs, Simon, and Fernando Leibovici. 2026. "Measuring Maritime Disruptions: A Route-Based Index of US Import Exposure." Federal Reserve Bank of Atlanta Policy Hub Paper 2026-9.

Fuchs, Simon, and Woan Foong Wong. 2026. "Multimodal transport networks." National Bureau of Economic Research.

Ganapati, Sharat, and Woan Foong Wong. 2023. "How far goods travel: global transport and supply chains from 1965–2020." *Journal of Economic Perspectives*, 37(3): 3–30.

Ganapati, Sharat, Woan Foong Wong, and Oren Ziv. 2024. "Entrepot: Hubs, scale, and trade costs." *American Economic Journal: Macroeconomics*, 16(4): 239–278.

Gaulier, Guillaume, and Soledad Zignago. 2010. "BACI: International Trade Database at the Product-Level: The 1994–2007 Version." CEPII CEPII Working Paper 2010-23.

Global Fishing Watch. 2025. "Global AIS Vessel Presence Dataset." *Data release*.

Grossman, Gene M., Elhanan Helpman, and Hugo Lhuillier. 2023. "Supply Chain Resilience: Should Policy Promote International Diversification or Reshoring?" *Journal of Political Economy*, 131(12): 3462–3496.

Hapag-Lloyd. 2023. "Update on the Red Sea and Suez Canal Situation." *Customer advisory*.

Hart, Brian, Matthew P. Funaiole, David Peng, Jasper Verschuur, Bonny Lin, and Leon Li. 2026. "Troubled Straits: Analyzing Trade Chokepoints in the South China Sea." Center for Strategic and International Studies.

Head, Keith, and Thierry Mayer. 2014. "Gravity equations: Workhorse, toolkit, and cookbook." In *Handbook of international economics*. Vol. 4, 131–195. Elsevier.

Heiland, Inga, Andreas Moxnes, Karen Helene Ulltveit-Moe, and Yuan Zi. 2025. "Trade from Space: Shipping Networks and the Global Implications of Local Shocks." *Review of Economics and Statistics*. doi:10.1162/rest.a.267.

Hirschman, Albert O. 1945. *National Power and the Structure of Foreign Trade.* Berkeley, CA:University of California Press.

Hulten, Charles R. 1978. "Growth Accounting with Intermediate Inputs." *Review of Economic Studies*, 45(3): 511–518.

Javorcik, Beata, Lucas Kitzmüller, Helena Schweiger, and Muhammed A. Yildirim. 2024. "Economic Costs of Friend-Shoring." *The World Economy*, 47(7): 2871–2908.

Jaworski, Taylor, Carl Kitchens, and Sergey Nigai. 2023. "Highways and globalization." *International Economic Review*, 64(4): 1615–1648.

Jiang, Bomin, Daniel Rigobon, and Roberto Rigobon. 2022. "From Just-in-Time, to Just-in-Case, to Just-in-Worst-Case: Simple Models of a Global Supply Chain under Uncertain Aggregate Shocks." *IMF Economic Review*, 70(1): 141–184.

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

Kalouptsidi, Myrto. 2018. "Detection and Impact of Industrial Subsidies: The Case of Chinese Shipbuilding." *Review of Economic Studies*, 85(2): 1111–1158.

Keohane, Robert O., and Joseph S. Nye. 1977. *Power and Interdependence: World Politics in Transition.* Boston, MA:Little, Brown and Company.

Kleinman, Benny, Ernest Liu, Stephen J. Redding, and David Xu. 2026. "International Trade in an Uncertain World." National Bureau of Economic Research NBER Working Paper 35416.

Lin, Bonny, Brian Hart, Matthew P. Funaiole, Samantha Lu, and Truly Tinsley. 2024*a*. "How China Could Blockade Taiwan." Center for Strategic and International Studies.

Lin, Bonny, Brian Hart, Matthew P. Funaiole, Samantha Lu, and Truly Tinsley. 2024*b*. "How China Could Quarantine Taiwan." Center for Strategic and International Studies.

Martin, Philippe, Thierry Mayer, and Mathias Thoenig. 2008. "Make Trade Not War?" *Review of Economic Studies*, 75(3): 865–900.

Maurer, Stephan, Luke Heath Milsom, and Ferdinand Rauch. 2026. "The Fragility of the Global Trading System." Centre for Economic Performance, London School of Economics Discussion Paper 2186.

Mediterranean Shipping Company. 2023. "MSC PALATIUM III Incident in Red Sea: Rerouting Suez Traffic to Cape." *Customer advisory*.

Milgrom, Paul, and Ilya Segal. 2002. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica*, 70(2): 583–601.

Mohr, Cathrin, and Christoph Trebesch. 2025. "Geoeconomics." Kiel Institute for the World Economy Kiel Working Paper 2279.

Notteboom, Theo, and Jean-Paul Rodrigue. 2008. "Containerisation, Box Logistics and Global Supply Chains: The Integration of Ports and Liner Shipping Networks." *Maritime Economics & Logistics*, 10(1–2): 152–174.

Palmer, Glenn, Roseanne W. McManus, Vito D'Orazio, Michael R. Kenwick, Mikaela Karstens, Chase Bloch, Nick Dietrich, Kayla Kahn, Kellan Ritter, and Michael J. Soules. 2022. "The MID5 Dataset, 2011–2014: Procedures, Coding Rules, and Description." *Conflict Management and Peace Science*, 39(4): 470–482.

Reyes-Heroles, Ricardo, and Xiaoqing Zhou. 2026. "Maritime Chokepoint Disruptions, Trade, and U.S. Inflation." Federal Reserve Bank of Dallas Working Paper.

Schelling, Thomas C. 1966. *Arms and Influence.* New Haven, CT:Yale University Press.

Seafarers' Rights International. 2018. "Cabotage Laws of the World." *Dataset and report*.

Simonovska, Ina, and Michael E. Waugh. 2014. "The Elasticity of Trade: Estimates and Evidence." *Journal of International Economics*, 92(1): 34–50.

S&P Global Market Intelligence. 2023–2024. "Panjiva Supply Chain Intelligence." *Shipment-level records of U.S. imports from bills of lading*, `https://panjiva.com`.

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

UNCTAD. 2023. "Review of Maritime Transport 2023: Towards a Green and Just Transition." United Nations Conference on Trade and Development UNCTAD/RMT/2023, Geneva.

UNCTAD. 2024. "Review of Maritime Transport 2024: Navigating Maritime Chokepoints." United Nations Conference on Trade and Development (UNCTAD) Report UNCTAD/RMT/2024, Geneva. eISSN: 2225-3459.

U.S. Department of Defense. 2024. "Military and Security Developments Involving the People's Republic of China 2024." U.S. Department of Defense.

USTR. 2025. "Notice of Action and Proposed Action in Section 301 Investigation of China's Targeting the Maritime, Logistics, and Shipbuilding Sectors for Dominance." Federal Register Notice dated April 17, 2025.

Verschuur, Jasper, Elco E. Koks, and Jim W. Hall. 2022. "Ports' Criticality in International Trade and Global Supply-Chains." *Nature Communications*, 13: 4351.

Verschuur, Jasper, Johannes Lumma, and Jim W. Hall. 2025. "Systemic Impacts of Disruptions at Maritime Chokepoints." *Nature Communications*, 16: 10421.

Vest, Charlie, Agatha Kratz, and Reva Goujon. 2022. "The Global Economic Disruptions from a Taiwan Conflict." Rhodium Group.

Wong, Woan Foong. 2022. "The round trip effect: Endogenous transport costs and international trade." *American Economic Journal: Applied Economics*, 14(4): 127–66.

A Analytical Results

This appendix proves the results of Section 2. The analytical benchmark holds productivity and amenities fixed and attributes no shipping income to households. Appendix A.5 explains how the quantitative model departs from it.

A.1 Competitive transport, Shephard's lemma, and traffic construction

This subsection derives the routing weight $\upsilon_{od,e}$ of equation (2.4) and its two interpretations. Carriers take edge costs as given and choose a feasible delivery plan $x \in \mathcal{R}_{od}$ to minimize the cost of delivering one unit, as in equation (2.3). Under additive costs, $C_{od}(x;\kappa) = \sum_e \kappa_e x_e$ is a transport bill. Under multiplicative costs, a plan is a path π and edge factors compound, $\tau_{od} = \min_{\pi \in \Pi_{od}} \prod_{e \in \pi} \kappa_e$. Taking logs gives an additive problem in traversal counts; Section 4 uses this specification.

Lemma A.1 (Shephard's lemma for routed transport)**.** *Let $\kappa \in \mathbb{R}^E_{++}$ and let $\mathcal{R}_{od}$ be a fixed, compact, nonempty feasible set. Assume C_{od} is continuously differentiable and nondecreasing in κ, and that C_{od} and its cost gradient are jointly continuous on the feasible set times an open neighborhood of κ. Write $\tau_{od}(\kappa) > 0$ for the minimum and $X^*_{od}(\kappa)$ for its minimizers.*

(i) (Danskin envelope.) τ_{od} has one-sided directional derivatives in any direction h,

$$\tau'_{od}(\kappa; h) = \min_{x \in X^*_{od}(\kappa)} \nabla_\kappa C_{od}(x;\kappa) \cdot h,$$

and is differentiable at κ whenever the gradient is common across optimal plans (in particular, when the optimal plan is unique), with

$$\frac{\partial \log \tau_{od}}{\partial \log \kappa_e} = \frac{\kappa_e}{\tau_{od}} \frac{\partial C_{od}(x^*_{od};\kappa)}{\partial \kappa_e} \equiv \upsilon_{od,e}. \tag{A.1}$$

(ii) (Additive case: cost shares.) At a differentiability point, if $C_{od}(x;\kappa) = \sum_e \kappa_e x_e$ with nonnegative service requirements x_e, then $\partial \tau_{od}/\partial \kappa_e = x^*_e$ and $\upsilon_{od,e} = \kappa_e x^*_e/\tau_{od} \in [0,1]$ is the cost share of edge e in delivering $o \to d$, with $\sum_e \upsilon_{od,e} = 1$ by degree-one homogeneity.

(iii) (Multiplicative case: usage intensities.) For path costs, $\upsilon_{od,e} = n_e(\pi^*)$ at a unique optimal path, where $n_e(\pi^*)$ is the number of traversals of edge e. It measures traversals rather than cost shares, so $\sum_e \upsilon_{od,e}$ equals the number of traversals on the chosen path. With Fréchet-distributed routing heterogeneity of shape θ, as in Section 4.2.1, the smoothed cost satisfies equation (4.7). For reachable pairs on a finite network whose continuation matrix has spectral radius below one, $\upsilon_{od,e} = \sum_\pi \omega_\pi n_e(\pi)$ is the expected traversal intensity. Path-choice probabilities ω_π weight traversal counts $n_e(\pi)$; the intensity is smooth and positive for edges on feasible $o \to d$ walks.

Proof. (i) General envelope. With a compact feasible set and a cost gradient that is jointly continuous in the plan and the edge costs, this is Danskin's theorem (see Milgrom and Segal, 2002). If all optimal gradients coincide, the value function is differentiable and the chain rule gives (A.1).

(ii) Additive costs. The objective gradient is x. Multiplying all edge costs by a positive scalar multiplies the objective and its minimum by that scalar; the feasible set is unchanged. Thus the value is homogeneous of degree one, and $\sum_e \kappa_e x^*_e = \tau_{od}$ gives the stated shares.

(iii) Path costs. On a finite feasible path set, apply (i) to log costs with $x_e = n_e(\pi)$. The same argument applies when positive cycle costs locally reduce the hard minimum to finitely many simple paths. For the smoothed walk index, let $w_\pi = \prod_{e \in \pi} \kappa_e^{-\theta}$ and $\omega_\pi = w_\pi / \sum_{\pi'} w_{\pi'}$. Then

$$\tau_{od}^{-\theta} = \sum_\pi w_\pi, \qquad \frac{\partial \log \tau_{od}}{\partial \log \kappa_e} = \sum_\pi \omega_\pi n_e(\pi).$$

This step differentiates the walk sum directly; part (i) does not apply because the set of walks is infinite. On a finite network, a spectral radius strictly below one persists in a neighborhood and permits differentiation of the convergent resolvent. The index and route weights are defined only for reachable origin–destination pairs. □

Two remarks. First, at a differentiability point the first-order response is evaluated at the baseline optimal plan. A tie between routes with different gradients requires the directional derivative in part (i); the smoothed recursion gives positive weights to feasible walks for each reachable pair. Second, the transport technology determines what exposure measures. Under multiplicative costs, $X_{od}\upsilon_{od,e}$ is shipment value times expected use of the edge. Under additive costs it is shipment value times the edge's share of the transport bill. If that bill is paid in addition to the goods price, the elasticity of the delivered price also includes the freight share of delivered value, the role of ξ_e in Section 4.

If scenario g changes generalized edge costs with traffic-specific equilibrium elasticities $\phi_{od,e,g}$, then

$$\frac{\partial \log \tau_{od}}{\partial d_g} = \sum_e \upsilon_{od,e}\phi_{od,e,g}. \tag{A.2}$$

Weighting by trade value, the direct effect on the delivered cost of trade from country i to country j is therefore

$$\begin{aligned}\Omega_{ij,g} &= \sum_{o\in\mathcal{N}_i}\sum_{d\in\mathcal{N}_j} X_{od}\frac{\partial \log \tau_{od}}{\partial d_g} \\ &= \sum_e \left[\sum_{o\in\mathcal{N}_i}\sum_{d\in\mathcal{N}_j} X_{od}\upsilon_{od,e}\phi_{od,e,g}\right] = \sum_e \Xi^X_{ij,e}\phi_{ij,e,g}.\end{aligned} \tag{A.3}$$

Here $\Xi^X_{ij,e}$ is the exposure defined in equation (2.6). Exposure is measured on baseline routed trade; the scenario enters through the equilibrium elasticities $\phi_{od,e,g}$. For common edge shocks, $\phi_{od,e,g} = \phi_{e,g}$ and equation (A.3) reduces to $\sum_e \Xi^X_{ij,e}\phi_{e,g}$. For selective shocks, $\phi_{ij,e,g}$ averages $\phi_{od,e,g}$ over location pairs, with weights proportional to each pair's contribution $X_{od}\upsilon_{od,e}$ to exposure.

A.2 Spatial welfare envelope and country aggregation

This subsection proves Proposition 2.1. Consider the Armington benchmark of Section 2, with fixed productivity and amenities, spending equal to labor income, and goods-market clearing. Assume that equilibrium outcomes, including edge costs and optimized routing costs, are differentiable at the baseline along scenario g. Country m's marginal equivalent-variation loss is

$$\mathcal{W}_{m,g} = \underbrace{\sum_{d\in\mathcal{N}_m}\sum_o X_{od}\frac{\partial \log \tau_{od}}{\partial d_g}}_{\text{delivery-cost effect}} + \underbrace{\sum_{d\in\mathcal{N}_m}\sum_o X_{od}r_{o,g} - \sum_{o\in\mathcal{N}_m} R_o r_{o,g}}_{\text{producer-price incidence}}, \tag{A.4}$$

where $r_{o,g}$ is the producer-price response at location o.

With amenities fixed, mobility gives $\mathrm{d}\log W_m = \mathrm{d}\log w_d - \mathrm{d}\log P_d$ at every populated location d of country m. Because welfare is equalized across a country's populated locations, the reallocation of workers adds no separate term. Weighting by baseline location income and summing gives

$$Y_m \mathrm{d}\log W_m = \sum_{d\in\mathcal{N}_m} y_d \,\mathrm{d}\log w_d - \sum_{d\in\mathcal{N}_m} y_d \,\mathrm{d}\log P_d, \qquad Y_m = \sum_{d\in\mathcal{N}_m} y_d. \tag{A.5}$$

Income weights allow direct substitution of goods-market clearing, $y_o = R_o = \sum_d X_{od}$, the fixed-productivity identity $\mathrm{d}\log w_o = \mathrm{d}\log p_o$, and the CES price index, $y_d \mathrm{d}\log P_d = \sum_o X_{od}(\mathrm{d}\log p_o + \mathrm{d}\log \tau_{od})$:

$$Y_m \mathrm{d}\log W_m = \sum_{o\in\mathcal{N}_m} R_o \mathrm{d}\log p_o - \sum_{d\in\mathcal{N}_m}\sum_o X_{od}\left(\mathrm{d}\log p_o + \mathrm{d}\log \tau_{od}\right). \tag{A.6}$$

The first sum values production at producer-price changes; the second values purchases at delivered-price changes.

Sales to domestic buyers appear in both and cancel. The producer-price terms therefore reduce to exports valued at domestic price changes minus imports valued at foreign price changes, a terms-of-trade term. The delivery-cost terms form the direct effect.

To construct the direct term, separate the producer-price terms, substitute the transport envelope, and change the order of summation. For a common edge shock,

$$\begin{aligned} -Y_m \left.\frac{\partial \log W_m}{\partial d_g}\right|_{r=0} &= \sum_{d\in\mathcal{N}_m}\sum_o X_{od}\frac{\partial \log \tau_{od}}{\partial d_g} \\ &= \sum_{d\in\mathcal{N}_m}\sum_o X_{od}\sum_e \upsilon_{od,e}\phi_{e,g} \\ &= \sum_e \left(\sum_{d\in\mathcal{N}_m}\sum_o X_{od}\upsilon_{od,e}\right)\phi_{e,g} = \sum_e \Xi^X_{m,e}\phi_{e,g} = \sum_i \Omega_{im,g}. \end{aligned} \tag{A.7}$$

Evaluating at $r = 0$ isolates the delivery-cost part of this differential; producer prices still adjust in equilibrium. Routing takes edge costs as given; the equilibrium of Section 4 (Definition 4.1) determines their response. Selective shocks use the pair-specific elasticities of equation (A.3). Substitution into equation (A.4) proves part (i).

For part (ii) of Proposition 2.1, suppose producer-price responses are common within each country. Define $\lambda_{ij} = X_{ij}/Y_j$ and, on $X_{ij} > 0$, $\varepsilon_{ij,g} = \Omega_{ij,g}/X_{ij}$; omit zero-flow pairs from ratios. Grouping (A.6) by country then gives

$$\ell_{j,g} = \sum_i \lambda_{ij}\big(\varepsilon_{ij,g} + r_{i,g} - r_{j,g}\big). \tag{A.8}$$

Multiplying by Y_j gives (2.10) with $\mathcal{T}^{\mathrm{ToT}}_{j,g} = \sum_i X_{ij} r_{i,g} - Y_j r_{j,g}$, since $\sum_i \lambda_{ij} = 1$ under the balanced-spending benchmark. This aggregation does not require fixed physical country output.

A.3 Recursive routing as a specialization

Recursive routing computes $\upsilon_{od,e}$ without changing the welfare formula. Use equation (4.7), with curvature $\theta = \sigma - 1 > 0$. Its unit term is a stopping option at the destination, but the continuation sum also permits walks that pass through the destination and return. On a finite network the sum is well defined when the continuation matrix has spectral radius below one; an acyclic network satisfies this condition automatically. Because it sums opportunities across walks, the inclusive cost is no larger than the cost of any one feasible path, and $\tau_{dd} \le 1$. Write $G_{ij} = \sum_{e:i\to j} \kappa_e^{-\theta}$ and $R = (I - G)^{-1}$, so $\tau_{od}^{-\theta} = R_{od}$. The identity $\mathrm{d}R = R(\mathrm{d}G)R$ implies, for $e = (i, j)$ and $R_{od} > 0$,

$$\upsilon_{od,e} = \frac{\partial \log \tau_{od}}{\partial \log \kappa_e} = \frac{R_{oi}\kappa_e^{-\theta}R_{jd}}{R_{od}}. \tag{A.9}$$

Under (4.5), bilateral expenditure is $X_{od} = p_o^{-\theta} R_{od} P_d^{\theta} y_d$. Multiplying by the route loading cancels R_{od} and gives

$$\begin{aligned} \sum_{o,d} X_{od}\upsilon_{od,e} &= \kappa_e^{-\theta}\left(\sum_o p_o^{-\theta} R_{oi}\right)\left(\sum_d R_{jd} P_d^{\theta} y_d\right) \\ &= \kappa_e^{-\theta}\mathcal{P}_i\mathcal{Q}_j = \Xi_e. \end{aligned} \tag{A.10}$$

The two bracketed sums satisfy the inward and outward recursions (4.8)–(4.9). Restricting the destination sum to country m gives its routed exposure $\Xi^X_{m,e}$. Thus (4.10) aggregates expenditure-weighted traversals rather than physical TEU. Arbitrary bilateral trade data need not have this factorization; the trade assignment of Section 3.1 instead assigns the observed country-pair values directly. Complete rotations and capacity restrictions affect the equilibrium edge-cost response $\phi_{e,g}$ through the supply system, which the routing envelope takes as given.

A.4 Proof of Proposition 2.2 and the disruptor's cost

This subsection proves Proposition 2.2 and relates its finite threshold to a marginal disruption. Let country a choose between a disruption g and leaving access open, O. Its payoff difference is

$$U_a(g;\chi) - U_a(O;\chi) = \chi L^{\text{fin}}_{h,g} - C^{\text{fin}}_{a,g}.$$

If the target loss $L^{\text{fin}}_{h,g}$ and the disruptor cost $C^{\text{fin}}_{a,g}$ are positive, disruption is weakly preferred if and only if $\chi \geq C^{\text{fin}}_{a,g}/L^{\text{fin}}_{h,g}$. This proves the finite comparison. If the target loss is positive but the disruptor cost is negative, disruption is strictly preferred for every $\chi \geq 0$. With zero disruptor cost it is weakly preferred, with indifference at $\chi = 0$.

For the marginal comparison, write the payoff at intensity d_g as

$$U_a(d_g;\chi) = V_a(d_g) + \chi Y_h[\log W_h(0) - \log W_h(d_g)]. \tag{A.11}$$

Using $\ell_{h,g} = -\partial \log W_h/\partial d_g$ and $C_{a,g} = -\partial V_a/\partial d_g$, the local derivative is

$$\left.\frac{\partial U_a(d_g;\chi)}{\partial d_g}\right|_{d_g=0} = \chi Y_h \ell_{h,g} - C_{a,g}. \tag{A.12}$$

Thus, if $Y_h\ell_{h,g} > 0$ and $C_{a,g} > 0$, the marginal payoff from increasing disruption intensity at the baseline is nonnegative if and only if

$$\chi \geq \chi^{\star}_{a,h,g} \equiv \frac{C_{a,g}}{Y_h \ell_{h,g}}. \tag{A.13}$$

For the benchmark in income units, set $V_a(\cdot) = Y_a \log W_a(\cdot)$, with Y_a fixed at baseline, and include no additional implementation cost. This normalization puts the disruptor's cost and the target's loss in the same units. In that case

$$C_{a,g} = \mathcal{W}_{a,g} = Y_a \ell_{a,g}, \qquad \chi^{\star}_{a,h,g} = \frac{\mathcal{W}_{a,g}}{\mathcal{W}_{h,g}}. \tag{A.14}$$

The threshold falls when the target's loss rises relative to the disruptor's. A broader implementation can add costs omitted from the disruptor's welfare,

$$C_{a,g} = \mathcal{W}_{a,g} + \Gamma^{\text{extra}}_{a,g} + \bar{c}_{a,g}. \tag{A.15}$$

Here $\Gamma^{\text{extra}}_{a,g}$ is a net transport-rent loss omitted from the reported welfare object, and $\bar{c}_{a,g}$ is a residual implementation cost. For a selective disruption, the transport-cost elasticities and welfare losses are computed under the selective technology.

Converting proportional losses to income units. The finite threshold compares losses in income units rather than proportional welfare losses. With the normalization of equation (A.14) and no additional rent or implementation cost, the disruptor cost is $C^{\text{fin}}_{a,g} = Y_a C_a(g)/100$ and the target loss is $L^{\text{fin}}_{h,g} = Y_h C_h(g)/100$, where $C_a(g)$ and $C_h(g)$ are welfare costs in log percentage points and Y_a and Y_h baseline incomes. Hence

$$\chi^{\text{fin}}_{a,h,g} = \frac{Y_a}{Y_h}\,\frac{C_a(g)}{C_h(g)}. \tag{A.16}$$

The income ratio Y_a/Y_h is therefore needed to convert a ratio of proportional losses into the threshold. The model does not assign capacity rentals to households (Section 4.1), so the shipping income a disruptor loses is outside $C_a(g)$ and would enter only through $\Gamma^{\text{extra}}_{a,g}$ in Equation (A.15), as would an implementation cost through $\bar{c}_{a,g}$.

A.5 Scope of the analytical benchmark

Write $\partial_g = \partial/\partial d_g$. Proposition 2.1 holds productivity and amenities fixed. In the quantitative model, $\widehat{A}_i = \widehat{L}_i^{\alpha}$ and $\widehat{u}_i = \widehat{L}_i^{\beta}$ add first-order local channels. To display them, use $\mathrm{d}\log w_i = \mathrm{d}\log p_i + \mathrm{d}\log A_i$ and $\mathrm{d}\log W_m =$

$\mathrm{d}\log u_i + \mathrm{d}\log w_i - \mathrm{d}\log P_i$ in the income-weighted aggregation. Writing $\mathcal{T}^{\mathrm{ToT}}_{m,g}$ for producer-price incidence in (A.4) gives

$$-Y_m \partial_g \log W_m = \sum_e \Xi^X_{m,e}\phi_{e,g} + \mathcal{T}^{\mathrm{ToT}}_{m,g} - \sum_{i\in\mathcal{N}_m} y_i(\partial_g \log A_i + \partial_g \log u_i). \tag{A.17}$$

With the exogenous components of productivity and amenities fixed, the final term becomes $-(\alpha+\beta)\sum_{i\in\mathcal{N}_m} y_i \partial_g \log L_i$. All weights are evaluated at the baseline. Because national labor is fixed, this term is zero to first order when each location's share of national labor equals its share of national income: then $\sum_{i\in\mathcal{N}_m} y_i\, \partial_g \log L_i$ is proportional to $\partial_g \sum_{i\in\mathcal{N}_m} L_i = 0$. The calibration assigns income and labor to a country's locations in the same proportions (Section 4.5). With spending weights, which apply when spending differs from income, this cancellation need not hold.

For the quantitative spending closure, let $\mu_i = E_i/y_i$ and $E_m = \sum_{i\in\mathcal{N}_m} E_i$. Mobility equalizes $W_m = u_i\mu_i w_i/P_i$, while CES shares are X_{od}/E_d. Multiplying its differential by baseline spending and summing within country gives

$$\begin{aligned} -E_m \partial_g \log W_m = &\sum_e \Xi^X_{m,e}\phi_{e,g} + \mathcal{T}^{\mathrm{ToT}}_{m,g} - \sum_{i\in\mathcal{N}_m} (E_i - y_i) r_{i,g} \\ &- \sum_{i\in\mathcal{N}_m} E_i(\partial_g \log A_i + \partial_g \log u_i + \partial_g \log \mu_i). \end{aligned} \tag{A.18}$$

All terms use the same baseline accounts. The additional producer-price term revalues the difference between spending and labor income, and the welfare denominator is now E_m rather than Y_m. A finite-change comparison must either impose the analytical benchmark or account separately for these channels.

B Equilibrium

This appendix states the equilibrium of Section 4 as a recursive system and in proportional changes.

B.1 Equilibrium in levels

Definition 4.1 in Section 4 defines the equilibrium in terms of agents' problems. Here we state it as a recursive system, substituting the optimal choices of households, producers, and operators into the market-clearing conditions. Two substitutions make this possible. Inward access is a power of the price index, $\mathcal{P}_i = P_i^{-\theta}$, because it sums the same power of the delivered prices of every variety that can reach i. Goods-market clearing then expresses local supply in access terms, $A_i^\theta w_i^{-\theta} = y_i/\mathcal{Q}_i$, so the bilateral cost indices τ_{ij} never have to be computed. Unlike Definition 4.1, spending can differ from labor income, so countries can run trade deficits. Under our spending rule, each location spends a fixed multiple of its income, scaled by a country factor that keeps the country's deficit at its baseline share of world income; with balanced trade, the two definitions coincide.

Definition B.1 (Equilibrium as a recursive system)**.** Given the network $(\mathcal{N}, \mathcal{E})$; fundamentals $\{\bar{A}_i, \bar{u}_i\}$ and labor endowments $\{\bar{L}_m\}$; exogenous edge costs $\{\bar{\kappa}_e\}$ and cargo values per unit of service $\{v_e\}$; port capacities $\{K^P_p\}$; pool capacities $\{K_h\}$; positive baseline spending–income ratios $\{\mu_i^0\}$ and country balances $\{\delta_m^0\}$ as shares of world income; and round-trip access $\{\Phi_{he}\}$, determined by operating requirements, base-port weights, and operational restrictions through (4.13), an equilibrium is a collection of wages $\{w_i\}$, populations $\{L_i\}$, welfare levels $\{W_m\}$, country spending multipliers $\{b_m\}$, inward and outward access $\{\mathcal{P}_i, \mathcal{Q}_i\}$, freight rates $\{q_e\}$, capacity rentals $\{\rho_h\}$, and generalized edge costs $\{\kappa_e\}$ such that, with labor income $y_i = w_i L_i$, spending $E_i = \mu_i y_i$ where $\mu_i = \mu_i^0 b_{m(i)}$, productivity $A_i = \bar{A}_i L_i^\alpha$, amenities $u_i = \bar{u}_i L_i^\beta$, and edge traffic $\Xi_e = \kappa_e^{-\theta}\mathcal{P}_{i(e)}\mathcal{Q}_{j(e)}$ for $e = (i(e), j(e))$, the following conditions hold:

(i) *Routing and market access.* For every location i,

$$\mathcal{P}_i = \frac{y_i}{\mathcal{Q}_i} + \sum_{k \in N^-(i)} \kappa_{ki}^{-\theta} \mathcal{P}_k, \qquad \mathcal{Q}_i = \frac{E_i}{\mathcal{P}_i} + \sum_{k \in N^+(i)} \kappa_{ik}^{-\theta} \mathcal{Q}_k. \tag{B.1}$$

(ii) *Goods markets.* For every location i, $y_i = A_i^{\theta} w_i^{-\theta} \mathcal{Q}_i$.

(iii) *Labor markets.* Welfare is equalized across the populated locations of each country, and populations sum to the country's labor force:

$$W_m = u_i \mu_i w_i \mathcal{P}_i^{1/\theta} \quad \text{for } i \in \mathcal{N}_m \text{ with } L_i > 0, \qquad \sum_{i \in \mathcal{N}_m} L_i = \bar{L}_m. \tag{B.2}$$

The multiplier b_m holds the country's nominal spending less labor income fixed as a share of world income:

$$\sum_{i \in \mathcal{N}_m} (E_i - y_i) = \delta_m^0 \sum_j y_j.$$

(iv) *Vessel capacity.* For every pool h, the rental is the value of its best uses, $\rho_h = \gamma_\zeta \big(\sum_{e \in \mathcal{E}} q_e^{\zeta} \Phi_{he}\big)^{1/\zeta}$.

(v) *Markets for shipping services.* For every edge e, demand equals supply:

$$\frac{\Xi_e}{v_e} = \gamma_\zeta^{\zeta} q_e^{\zeta - 1} \sum_h \Phi_{he} \rho_h^{1-\zeta} K_h. \tag{B.3}$$

(vi) *Generalized edge costs.* For every edge e, freight payments and port congestion feed back into costs:

$$\kappa_e = \bar{\kappa}_e \Big(1 + \frac{q_e}{v_e}\Big) \prod_{p \in \{i(e), j(e)\}} \Big(\frac{1}{K_p^P} \sum_{e' \in \mathcal{E}_p} \iota_{pe'} \Xi_{e'}\Big)^{\lambda_p}, \tag{B.4}$$

where $\mathcal{E}_p$ is the set of edges that begin or end at port p and ι_{pe} the handling required at p per unit of traffic on e.

Wages are determined up to the choice of numeraire; we normalize world income, $\sum_i y_i$, to one.

B.2 Baseline shares and proportional changes

Baseline shares. In the level model, B_{he}^0 is baseline service revenue generated by pool h on edge e. The implementation uses a deployment- and distance-weighted revenue proxy. Define

$$s_{he}^0 = \frac{B_{he}^0}{\sum_{e' \in \mathcal{E}} B_{he'}^0}, \qquad \widetilde{s}_{he}^0 = \frac{B_{he}^0}{\sum_{h'} B_{h'e}^0}. \tag{B.5}$$

The first object is the share of pool h's freight revenue earned on edge e; the second is pool h's share of the revenue earned on that edge. Here h denotes an operator-country pool, not a base port. Because all pools on an edge receive the same freight rate, $\widetilde{s}_{he}^0$ is also the pool's service-supply share. The rental equation uses s_{he}^0 to value a pool's opportunities, while the supply equation uses $\widetilde{s}_{he}^0$ to aggregate services across pools.

The full system. Dividing every condition of Definition B.1 by its baseline value (Section 4.4) leaves three sets of baseline shares. The first describes routing. The throughput of a location, $\mathcal{P}_i \mathcal{Q}_i$, equals its income plus the value arriving by sea, and also its spending plus the value leaving; the routing shares split it into its sources and uses,

$$\omega_{ii}^P = \frac{y_i^0}{\mathcal{P}_i^0 \mathcal{Q}_i^0}, \quad \omega_e^P = \frac{\Xi_e^0}{\mathcal{P}_{j(e)}^0 \mathcal{Q}_{j(e)}^0}, \qquad \omega_{ii}^Q = \frac{E_i^0}{\mathcal{P}_i^0 \mathcal{Q}_i^0}, \quad \omega_e^Q = \frac{\Xi_e^0}{\mathcal{P}_{i(e)}^0 \mathcal{Q}_{i(e)}^0},$$

where y_i^0 and E_i^0 are baseline labor income and spending, so that inward shares sum to one over local production and arriving edges and outward shares over local spending and departing edges. The second set is the revenue

shares s^0_{he} and $\widetilde{s}^0_{he}$ of equation (B.5), which combine the capacity of the schedules with the traffic of the trade assignment (Section 3.1). Capacity determines who supplies an edge, and traffic determines how much the edge is worth and how much shipping it demands. The third describes costs: ξ_e, the share of the freight payment in the edge cost, and ς^0_{pe}, edge e's share of port p's throughput.

Proposition B.1 (Counterfactual equilibrium in proportional changes)**.** *Given the baseline shares, the parameters $(\theta, \zeta, \alpha, \beta, \{\lambda_p\})$, and changes in fundamentals $\widehat{\widetilde{A}}_i$ and $\widehat{\widetilde{u}}_i$, exogenous edge costs $\widehat{\widetilde{\kappa}}_e$, cargo values $\widehat{v}_e$, port capacities $\widehat{K}^P_p$, pool capacities $\widehat{K}_h$, and round-trip access $\widehat{\Phi}_{he}$, the counterfactual equilibrium is a solution $\{\widehat{\mathcal{P}}_i, \widehat{\mathcal{Q}}_i, \widehat{w}_i, \widehat{L}_i, \widehat{W}_m, \widehat{q}_e, \widehat{\rho}_h, \widehat{\kappa}_e\}$ to the following system, in which $\widehat{y}_i = \widehat{w}_i\widehat{L}_i$ is the change in location i's income, $\widehat{E}_i = \widehat{\mu}_i\widehat{y}_i$ the change in its spending, with $\widehat{\mu}_i$ the change in its ratio of spending to income, common to the locations of a country, that holds the country's nominal balance fixed in units of world income, $\widehat{A}_i = \widehat{\widetilde{A}}_i\widehat{L}_i^{\alpha}$, $\widehat{u}_i = \widehat{\widetilde{u}}_i\widehat{L}_i^{\beta}$, and $\widehat{\Xi}_e = \widehat{\kappa}_e^{-\theta}\widehat{\mathcal{P}}_{i(e)}\widehat{\mathcal{Q}}_{j(e)}$:*

(i) Routing and market access:

$$\widehat{\mathcal{P}}_i = \omega^P_{ii}\frac{\widehat{y}_i}{\widehat{\mathcal{Q}}_i} + \sum_{e:\, j(e)=i} \omega^P_e\, \widehat{\kappa}_e^{-\theta}\widehat{\mathcal{P}}_{i(e)}, \qquad \widehat{\mathcal{Q}}_i = \omega^Q_{ii}\frac{\widehat{E}_i}{\widehat{\mathcal{P}}_i} + \sum_{e:\, i(e)=i} \omega^Q_e\, \widehat{\kappa}_e^{-\theta}\widehat{\mathcal{Q}}_{j(e)}.$$

(ii) Goods and labor markets: $\widehat{y}_i = \widehat{A}_i^{\theta}\widehat{w}_i^{-\theta}\widehat{\mathcal{Q}}_i$, $\widehat{W}_{m(i)} = \widehat{\mu}_i\widehat{u}_i\widehat{w}_i\widehat{\mathcal{P}}_i^{1/\theta}$, $\sum_{i\in\mathcal{N}_m} L_i^0\widehat{L}_i = \bar{L}_m$, *and* $\sum_{i\in\mathcal{N}_m}(E_i^0\widehat{E}_i - y_i^0\widehat{y}_i) = \sum_{i\in\mathcal{N}_m}(E_i^0 - y_i^0)$.

(iii) Vessel capacity: $\widehat{\rho}_h^{\zeta} = \sum_e s^0_{he}\,\widehat{q}_e^{\,\zeta}\,\widehat{\Phi}_{he}$.

(iv) Markets for shipping services: $\widehat{\Xi}_e/\widehat{v}_e = \widehat{q}_e^{\,\zeta-1}\sum_h \widehat{s}^0_{he}\,\widehat{\Phi}_{he}\,\widehat{\rho}_h^{1-\zeta}\widehat{K}_h$.

(v) Generalized edge costs:

$$\widehat{\kappa}_e = \widehat{\widetilde{\kappa}}_e\Big[(1-\xi_e) + \xi_e\frac{\widehat{q}_e}{\widehat{v}_e}\Big]\prod_{p\in\{i(e),j(e)\}}\Big(\frac{1}{\widehat{K}^P_p}\sum_{e'}\varsigma^0_{pe'}\widehat{\Xi}_{e'}\Big)^{\lambda_p}.$$

We normalize world income by $\sum_i \varpi_i\widehat{y}_i = 1$, with ϖ_i location i's baseline share of world income.

Derivation: divide the conditions of Definition B.1 by their baseline values. The spending rule gives $\widehat{\mu}_i = b_{m(i)}$, with baseline $b^0_m = 1$; subtracting the normalized baseline balance yields condition (ii). Appendix B.3 gives the remaining substitutions for conditions (iii) and (iv).

B.3 Complete rotations in levels and proportional changes

Round-trip access Φ_{he} combines the access indices of equation (4.12) as in equation (4.13). For each movement type, we require the matrix of arc weights $(g^X_{ik})^{-\zeta}$ to have spectral radius below one, which keeps the sum in (4.12) finite. The rental and supply equations in levels are

$$\rho_h^{\zeta} = \gamma_{\zeta}^{\zeta}\sum_e q_e^{\zeta}\Phi_{he}, \tag{B.6}$$

$$S^S_e = \gamma_{\zeta}^{\zeta} q_e^{\zeta-1}\sum_h \Phi_{he}\rho_h^{1-\zeta}K_h. \tag{B.7}$$

To see where these equations come from, let Z_{he} denote effective service productivity for capacity pool h on a complete rotation serving edge e, with independent draws across eligible edges and $\Pr(Z_{he} \le z) = \exp\{-\Phi_{he}z^{-\zeta}\}$. Assume $\zeta > 1$, a finite positive opportunity sum, and $\gamma_\zeta = \Gamma(1-1/\zeta)$. Capacity is contracted before the draws, or diversified across a continuum of draws, so the common rental equals expected maximum revenue. A capacity unit is assigned to the edge that maximizes $q_e Z_{he}$. Its choice probability is

$$\pi_{he} = \frac{q_e^{\zeta}\Phi_{he}}{\sum_a q_a^{\zeta}\Phi_{ha}}, \tag{B.8}$$

and expected maximum revenue is the rental index in (B.6), since independence gives $\Pr(\max_e q_e Z_{he} \le z) = \exp\{-z^{-\zeta} \sum_e q_e^{\zeta} \Phi_{he}\}$ for $z > 0$. The common rental is an ex ante or diversified zero-profit condition; realized revenue need not equal the rental on each capacity unit. The envelope theorem then gives the amount of service supplied by pool h on edge e as $K_h \partial \rho_h / \partial q_e$. Summing over pools yields (B.7). Equivalently, $S_{he} = K_h \rho_h \pi_{he} / q_e$: revenue $K_h \rho_h$ is allocated across services by π_{he}, and division by the freight rate converts revenue into service units. This quantity includes the selection of the capacity units best suited to that service.

Dividing baseline and counterfactual equations gives the rental and supply conditions (iii) and (iv) of Proposition B.1. Multiplying the level supply equation by q_e and summing across edges, then using the rental equation, yields

$$\sum_e q_e S_e^S = \sum_h \rho_h K_h. \tag{B.9}$$

Freight revenue is thus paid out in full as capacity rentals. This level identity follows from the degree-one homogeneity of each pool's rental index in freight rates.

Baseline requirements and base-port weights. In the baseline, the operating requirement of edge $e = (i, j)$ is $g_e = (\omega_e^Q)^{-1/\zeta}$, where ω_e^Q is its flow divided by income plus incoming flows at its tail node. Thus $(g_e)^{-\zeta} = \omega_e^Q$ for every ζ, and the same arc weights enter the ballast, laden, and return recursions. Write $G_{ij} = \sum_{e:i\to j} g_e^{-\zeta}$ and $\mathcal{R} = (I - G)^{-1}$. With a fixed service wedge a_{he}^0, equation (4.13) gives

$$\Phi_{he}^0 = (a_{he}^0)^{-\zeta} \sum_n b_{hn} \mathcal{R}_{ni} \mathcal{R}_{ij} \mathcal{R}_{jn}.$$

We set $a_{he}^0 = 1$; a fixed wedge cancels from $\widehat{\Phi}_{he}$.

A closure multiplies affected operating requirements by their route-distance ratios. The accompanying change in the exogenous edge cost sets $\widehat{\bar{\kappa}}_e$ to the square root of edge e's route-distance ratio, as in Sections 5.1 and 6.1. Base-port weights b_{hn} allocate half of each pool's baseline revenue proxy on an edge to each endpoint, normalize within pool, and remain fixed in counterfactuals.