March 2026

Geopolitical Risk and Foreign Dependence in Global Shipping[1]

Simon Fuchs
Federal Reserve Bank of Atlanta

Fernando Leibovici
Federal Reserve Bank of St. Louis

Woan Foong Wong
University of Oregon
CEPR & NBER

Abstract

We study countries' dependence on foreign-built and foreign-operated containerships, the vulnerability this dependence creates to shipping and geopolitical disruptions, and the potential for policy to mitigate these risks. We document that global shipbuilding is highly concentrated in China and South Korea, with South Korea specializing in larger vessels and China accounting for a growing share of newer capacity. In contrast, vessel operation is substantially more dispersed: most Chinese-built ships are operated by non-Chinese carriers, highlighting the deep integration of Chinese-built capacity into global shipping networks. Using ship-level registry data matched to the universe of U.S. import shipments, we show that the United States mirrors this global structure: U.S. trade relies overwhelmingly on vessels from East Asia while the operators serving U.S. routes are considerably more diversified. To quantify how such dependence shapes trade costs and vulnerability, we develop a quantitative shipping model with endogenous fleet investment, rental markets, route assembly by operators, and spatial general equilibrium trade. The framework provides a tractable mapping from fleet-specific shocks, whether arising from geopolitical disruptions or policy interventions, to shipping prices, routing patterns, bilateral trade costs, and aggregate welfare. By combining observed exposure to different fleets with shipping cost shares and trade patterns, the model allows us to evaluate how disruptions propagate and how policies that target particular fleets reshape trade outcomes in equilibrium. Applied to recent Section 301-style surcharges, the model implies welfare losses concentrated on highly exposed routes, with additional spillovers operating through fleet substitution and routing adjustment.

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

[1]Contact: simon.fuchs@atl.frb.org, fernando.leibovici@stls.frb.org, and wfwong@uoregon.edu. We thank Myrto Kalouptsidi for an insightful discussion and the participants at the 2025 NBER Conference on Transportation Networks and the Spatial Distribution of Economic Activity for their valuable feedback. We also appreciate the helpful feedback from seminar participants at Monash University and University of Melbourne. KC Pringle and Valerie Nguyen provided excellent research assistance. The views expressed herein are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve Bank of Atlanta, the Federal Reserve System, or the Board of Governors. All errors are our own.

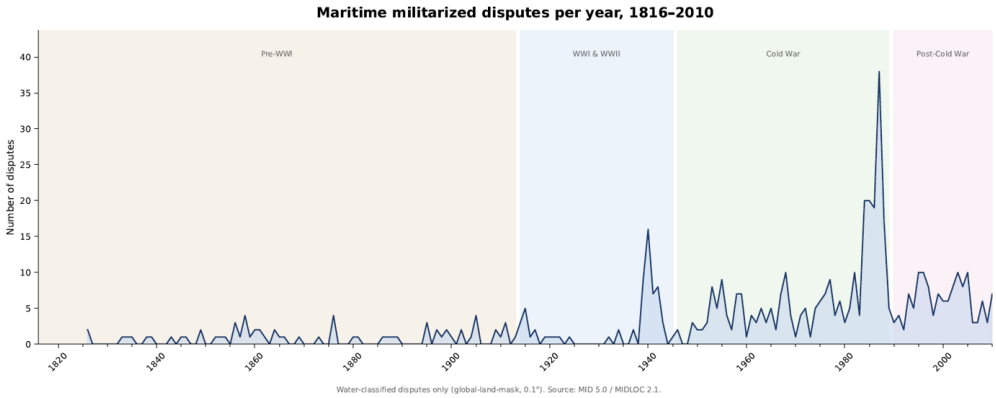
1 Introduction

Global maritime shipping accounts for more than 80 percent of world trade, and the cost and availability of shipping services depend critically on the structure of the global containership fleet. Yet the seas over which trade flows have never been free from conflict. Figure 1a documents the frequency of maritime militarized disputes from 1816 to 2010, revealing that such confrontations have intensified markedly since the mid-twentieth century, with a sharp spike during the Cold War era and continued elevated levels in the post-Cold War period. Figure 1b maps the geographic distribution of these disputes and shows that they cluster disproportionately near major maritime chokepoints, like the Strait of Hormuz, the Strait of Malacca, the Suez Canal, the Bab el-Mandeb, and other critical transit corridors through which the bulk of global trade passes. These patterns underscore that the physical infrastructure of global trade is embedded in a geopolitically contested environment, where access to shipping lanes can be disrupted by state conflict, blockades, or militarized confrontations.

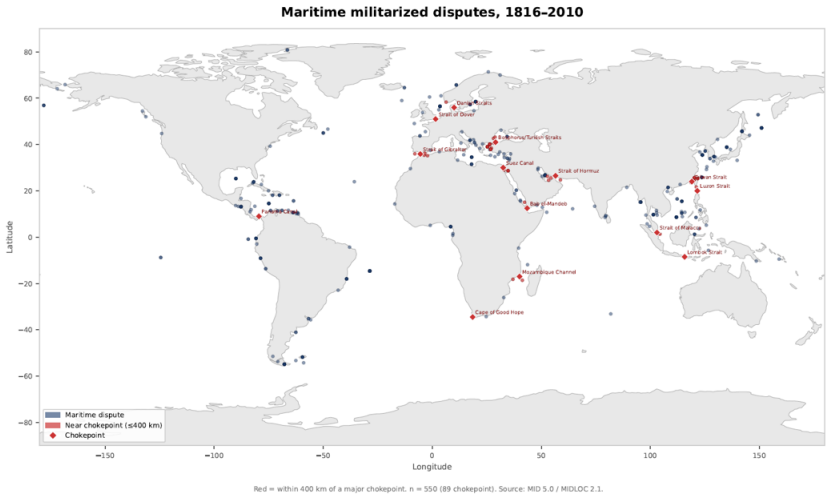
Governments have long recognized this vulnerability. Figure 2 maps the global prevalence of cabotage laws, which reserve domestic shipping routes for nationally flagged or domestically owned vessels. These restrictions are neither recent nor confined to a handful of countries: they span major trading economies on every continent, from the United States and Australia to Brazil, China, and much of Europe. Cabotage laws represent the most direct policy expression of a concern that is far broader in scope. Across the world, governments impose a wide array of barriers to foreign participation in shipping markets, including restrictions on foreign operators, cargo reservation schemes, mandates to use domestic port services, and privileges for state-owned carriers (OECD STRI, 2025). These patterns reflect a long-standing and widespread judgment that reliance on foreign-controlled shipping capacity constitutes a strategic vulnerability, one that warrants active policy intervention.

Against this backdrop, the structure of the global shipbuilding industry has undergone a profound transformation. Historically, the United States, United Kingdom, and other European countries maintained substantial shipbuilding industries and operated large commercial fleets (Barwick, Kalouptsidi and Zahur, 2024, 2025). Today, by contrast, these countries play a much more limited role in international shipping. Over the same period, shipbuilding has become highly concentrated in East Asia, where sustained industrial policies and large-scale production have enabled China, South Korea, and Japan to dominate new deliveries. This shift concentrates critical transport capacity in a small number of foreign economies, creating exposure not only to commercial shipping disruptions but also to the broader geopolitical tensions that increasingly characterize relations among major powers.

In this paper, we study whether such dependence represents a source of vulnerability that warrants policy intervention, and whether interventions that alter fleet composition strengthen resilience or instead primarily raise trade costs. We assemble new microdata linking the global registry of containerships to shipment-level trade flows, which allows us to measure exposure to foreign fleets at a granular level across countries, routes, vessel sizes, and vintages. For the United States, we match these vessels to the universe of import shipments, providing a detailed view of how global fleet concentration translates into shipment-level exposure. To interpret these patterns and evaluate the effects of fleet-specific disruptions or policy interventions, we develop a tractable framework that embeds an explicit shipping sector within a spatial general equilibrium model of trade. The model connects changes in fleet composition and



(a) Maritime militarized disputes per year, 1816–2010



(b) Geographic distribution of maritime militarized disputes

Figure 1. Maritime Militarized Disputes, 1816–2010. Panel (a) reports the annual count of water-classified militarized interstate disputes. Background shading indicates major geopolitical eras. Panel (b) maps their geographic distribution. Blue circles indicate dispute locations; red shading marks disputes within 400 km of a major chokepoint (red diamonds). Source: MID 5.0 / MIDLOC 2.1.

fleet-specific costs to shipping prices, routing patterns, bilateral trade costs, and aggregate welfare. By combining observed exposure to different fleets with shipping cost shares and trade patterns, the framework allows us to evaluate how disruptions propagate through shipping markets and how policies that target particular fleets reshape trade outcomes in equilibrium.

Our analysis yields four main findings. First, global shipbuilding is highly concentrated in East Asia, and countries such as the United States rely predominantly on foreign-built vessels drawn from this concentrated supply base, with China's presence especially pronounced among newer and smaller ships. Second, this concentration reflects sizable efficiency advantages: unit prices decline sharply with vessel size and are generally lower for China-built ships, indicating that access to large foreign shipyards delivers important efficiency gains. Third, the same concentration creates measurable vulnerability. For instance, when access to Chinese-built capacity is restricted in the model, shipping prices rise disproportionately on routes that rely heavily on these fleets, and the resulting trade cost increases propagate through rental markets and routing adjustments across the network. The magnitude of these effects depends on

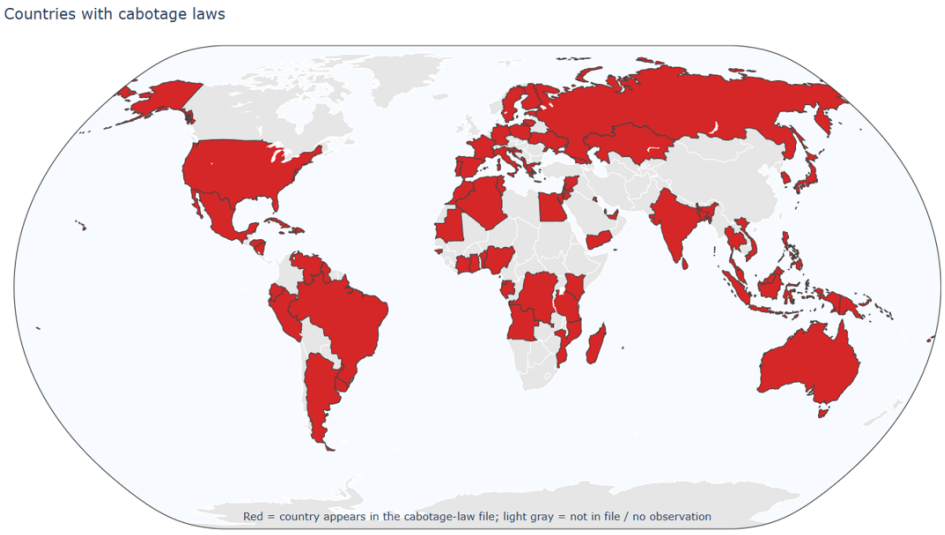


Figure 2. Countries with Cabotage Laws. Red shading indicates countries with cabotage laws that reserve domestic shipping routes for nationally flagged or domestically owned vessels. Light gray indicates countries not covered or without such restrictions. Cabotage restrictions are widespread across both developed and developing economies, spanning every continent. Source: Seafarers' Rights International.

route-level exposure, substitution possibilities across fleets, and the share of shipping in delivered prices. Fourth, policies that reduce reliance on Chinese-linked fleets lower vulnerability to such disruptions but do so at a cost. Applied to the proposed Section 301 surcharges, the model implies welfare losses concentrated on highly exposed lanes, reflecting the trade-off between reduced dependence and higher trade costs.

Taken together, the data and the framework characterize how concentration in the global shipping fleet shapes both the efficiency and the vulnerability of international trade to shipping and geopolitical disruptions. They show how reliance on foreign-built capacity influences the cost and reliability of trade, how disruptions to specific fleets propagate through shipping markets and trade networks, and how exposure varies systematically across routes and countries. More broadly, the analysis highlights how control over upstream transportation infrastructure can transmit geopolitical and economic shocks across borders, and it provides a disciplined framework for evaluating policies that seek to reduce dependence while weighing the associated efficiency costs.

In the first part of the paper, we assemble and analyze new data on the composition of the global containership fleet and its role in international trade. We combine vessel-level information from Clarksons Shipping Intelligence Network with shipment-level U.S. import records from Panjiva, matching ships to individual port calls and import flows. This linked dataset allows us to measure, at the level of routes, ship sizes, and vintages, how much capacity is built in each country and operated by carriers of each nationality, and to compare the structure of fleets serving U.S. trade to the global distribution of shipbuilding.

We use these data to document a set of stylized facts about exposure and concentration. We show that global shipbuilding is heavily concentrated in East Asia, with South Korea and China accounting for most of the world's containership capacity across size classes and vessel ages. U.S. imports reflect this concentration, since most ships serving U.S. trade are foreign built and foreign operated, and East

Asian yards supply the fleets that dominate major U.S. trade lanes. China's role has increased noticeably in more recent cohorts, particularly among newer and smaller vessels, highlighting the ongoing shifts within this concentrated global supply base.

We then examine the cost side of this dependence by studying prices for ships in the current containership orderbook. Using Clarksons data on contract prices and capacities, we document strong scale economies in shipbuilding, with prices per unit of capacity declining sharply as vessel size increases, and we show that China-built ships are generally cheaper than comparable vessels built elsewhere. These patterns indicate that access to large foreign shipyards delivers substantial cost advantages, and they provide the empirical basis for calibrating the model's shipping cost structure and for evaluating the incidence of policies that target particular fleets.

Finally, we distinguish between the countries that build vessels and the firms that operate them. Although shipbuilding is heavily concentrated in China, South Korea, and Japan, the operation of containerships is much more internationally dispersed. Ships constructed in these East Asian yards are frequently deployed by carriers headquartered in Europe, the United States, and other parts of Asia, and the fleets serving U.S. trade reflect this broad mix of operator nationalities. This separation between builder and operator origin indicates that exposure through ship construction and exposure through the deployment of vessels are distinct margins. It also underscores that modern containerships are embedded in a global network in which construction, ownership, and operation span multiple countries.

The second part of the paper develops a model that connects the structure of the global containership fleet to shipping prices, trade costs, and welfare in a stochastic shipping environment. The model captures three central features of maritime transport. First, shipping services are assembled from heterogeneous fleets that are owned by different countries and deployed globally across trade routes. Second, the cost of moving goods depends on the composition of the fleets serving each route and on the routing decisions that arise in spatial equilibrium. Third, both fleet availability and route-level shipping conditions are subject to uncertainty, which affects rental rates, routing choices, and ultimately delivered trade costs. These elements allow us to study how shocks to specific fleets or to the cost of using particular fleets propagate through the shipping network and shape trade flows and welfare.

The framework embeds an explicit shipping sector within a spatial general equilibrium model of trade. Fleet producers choose capacity ex-ante, before uncertainty is realized about where their vessels will be able to operate. Ex-post, in each state of the world, operators rent whatever capacity is available and deploy it across routes, facing congestion and competing with vessels from other fleet origins. Shippers choose among alternative routes based on their relative costs and reliability, directing traffic toward paths that are cheaper or more dependable and away from those exposed to greater uncertainty. Delivered trade costs therefore reflect the mix of fleets active on each link, the state-contingent rental rates that emerge in equilibrium, and the adjustments in routing that reallocate traffic across the network. Together, these elements provide a tractable way to study how shocks to specific fleets propagate through shipping markets and affect bilateral trade costs and welfare.

The model delivers a clear mapping from fleet specific shocks to shipping prices, trade costs, and welfare. A change in the cost or availability of a particular fleet affects the routes that rely on it directly, and it also influences shipping conditions on other routes through the rental market for capacity and through adjustments in routing. These propagation channels reflect the structure of the global shipping

network. Direct exposure determines which routes are immediately affected, substitution across fleets and corridors shapes how operators respond, and the network position of each route determines how shocks transmit through the system. These mechanisms characterize how the organization of shipping capacity influences delivered trade costs and the allocation of international shipments.

In the third part of the paper, we use the framework to quantify the vulnerability implied by concentration in Chinese-built capacity. We consider a counterfactual in which access to Chinese-built vessels on U.S.-serving routes is reduced. In the model, such a disruption raises shipping costs disproportionately on routes that rely heavily on these fleets, with effects that propagate through rental markets and routing adjustments across the network. The resulting trade cost increases depend on route-level exposure, substitution possibilities across fleets, and the importance of shipping in delivered prices. This exercise provides a quantitative measure of how concentration in shipbuilding translates into vulnerability in trade outcomes.

We then evaluate policies designed to mitigate such exposure. We study the surcharge package proposed by the United States in 2025 for vessels built in Chinese shipyards or operated by Chinese firms. These measures would have raised the effective cost of deploying Chinese-linked fleets on U.S. routes. We convert the statutory fee schedule into ad valorem cost increases by vessel size class and treat these increases as fleet-specific shocks within the model. This mapping allows the empirical patterns of exposure to interact directly with the model's propagation structure and provides a transparent way to assess how policies that reduce reliance on particular fleets affect shipping prices, trade costs, and welfare.

Finally, we aggregate these route-level effects to quantify their macroeconomic implications and to clarify the underlying trade-off. The resulting changes in delivered prices mirror underlying exposure: lanes that rely intensively on Chinese-linked vessels experience the largest increases in shipping costs, while routes served primarily by other fleets are less affected. Aggregating across routes implies welfare losses on the order of 0.01 to 0.02 percent of GDP under the proposed surcharge schedule. These results highlight the central tension identified in the paper: policies that reduce reliance on concentrated foreign fleets lower exposure to future disruptions but do so by raising trade costs and reducing the efficiency gains associated with access to large, low-cost shipyards. The magnitude of this trade-off reflects direct exposure, substitution possibilities across fleets and corridors, and the relatively modest share of shipping costs in delivered import prices.

The results underscore the importance of understanding how the organization of global shipping interacts with domestic trade outcomes. The evidence reveals a concentrated foreign supply base that is central to the global movement of goods, and the model shows how cost changes affecting particular fleets propagate through shipping prices, trade flows, and welfare. By combining detailed microdata with a tractable general equilibrium framework, the paper provides an approach for analyzing how upstream transport technologies, produced and operated abroad, shape the effective cost of market access. This perspective is essential for evaluating policies that seek to modify the composition of the fleets that serve U.S. trade and for assessing the broader economic consequences of reliance on foreign shipping capacity.

This paper connects to the literature studying how transport networks shape trade flows and welfare in spatial equilibrium. Work in this area shows how link-level costs propagate through networks and influence aggregate outcomes, including Allen and Arkolakis (2014, 2022), Fajgelbaum and Schaal

(2020), Fan and Luo (2020), Fan, Lu and Luo (2023), Jaworski, Kitchens and Nigai (2023), Bonadio (2021) and Fuchs and Wong (2024). Related contributions examine congestion, routing, and market access, such as Coşar and Fajgelbaum (2016); Coşar and Demir (2016), Tsivanidis (2022), Kreindler and Miyauchi (2023), and Miyauchi, Nakajima and Redding (2025). Our contribution is to embed foreign-owned shipping supply directly into this network structure and to characterize the general equilibrium incidence of shocks that affect specific fleets.

The paper also relates to research on shipping markets and maritime trade. This literature highlights the role of endogenous shipping supply, operating decisions, and fleet heterogeneity in determining trade costs and shaping global trade patterns. Key contributions include Kalouptsidi (2014), Brancaccio, Kalouptsidi and Papageorgiou (2020), Barwick, Kalouptsidi and Zahur (2024, 2025), Heiland et al. (2019), Ganapati, Wong and Ziv (2024), Dunn and Leibovici (2024), Ganapati and Wong (2023) and Wong (2022). Empirical work on shipping networks and transport costs, such as Brooks, Gendron-Carrier and Rua (2018), Notteboom and Rodrigue (2008); Rodrigue, Notteboom and Pallis (2011), Cristea et al. (2013), Shapiro (2016), and Lugovskyy, Skiba and Terner (2025), documents the centrality of maritime transport in global trade. Our analysis contributes by linking vessel-level data to U.S. import flows and incorporating a detailed shipping sector into a quantitative trade model.

Finally, the paper speaks to the growing literature on geoeconomics and strategic dependence in global production and infrastructure. This work examines how reliance on foreign-controlled inputs shapes vulnerability, policy design, and the transmission of shocks, including Farrell and Newman (2019, 2023), Blackwill and Harris (2016), Baldwin (2022), Rodrik (2022), Rodrik and Sabel (2019), and Antràs (2020); Antràs and Chor (2022). Related contributions study exposure in trade networks and the macroeconomic implications of foreign dependence, such as Bonadio et al. (2021), Freund et al. (2020), Grossman, Helpman and Lhuillier (2023), and Gawande, Krishna and Olarreaga (2009). Our paper adds to this literature by quantifying U.S. dependence on foreign-controlled shipping capacity and by developing a framework to analyze how this dependence affects trade costs and welfare.

The remainder of the paper is organized as follows. Section 2 presents the data and documents the empirical patterns that characterize the global distribution of containerships and the exposure of U.S. trade to foreign fleets. Section 3 develops the model, describes its main mechanisms, and derives the relationships that link fleet composition to shipping prices and trade costs. Section 4 applies the model to the Section 301 fees and presents the quantitative results on shipping prices, trade flows, and welfare. Section 5 discusses the quantitative implementation and calibration. Section 6 concludes.

2 A Simple Model of Trade and Infrastructure

This section develops an extension of the (Dornbusch, Fischer and Samuelson, 1977) model to study how dependence on foreign-owned or foreign-controlled shipping capacity shapes delivered comparative advantage, trade patterns, and the value of capacity policy. The trade side remains DFS in spirit: equilibrium is still governed by a monotone comparison of delivered costs across a continuum of goods. The key departure from the benchmark is that delivered trade costs are no longer primitive. Instead, they are generated by an endogenous routing problem in which imported goods are allocated across

a continuum of shipping locations subject to congestion. This structure is meant to capture a central concern in the paper: when transport capacity is controlled abroad, geopolitical manipulation can translate into strategic disruptions of shipping services and hence into higher delivered costs for the goods that rely on them. We model that vulnerability in reduced form through stochastic usable capacity. The section proceeds in three steps. It first characterizes a risk-free benchmark in which the transport wedge is determined solely by routing and congestion. It then introduces risk by allowing usable capacity to become state-dependent, which generates higher expected transport costs and a reliability motive for policy. Finally, it derives sharp implications for the optimal ex-ante policy response, first in terms of broad macro elasticities and then, in the analytically tractable benchmark, in closed form. Detailed derivations are collected in Appendix A.3–A.8, so the main text focuses on economic interpretation.

Setup. There are two countries, Home and Foreign, one factor of production, labor, and a continuum of goods indexed by $g \in [0,1]$. Home consumers have log-separable preferences, so if $b(g)$ denotes the expenditure weight on good g, with $b(g) \geq 0$ and $\int_0^1 b(g)\,\mathrm{d}g = 1$, indirect utility takes the form $V_H = \ln E_H - \int_0^1 b(g) \ln p_H(g)\,\mathrm{d}g$, where E_H is Home expenditure and $p_H(g)$ is the delivered Home-market price of good g. On the production side, unit labor requirements are $a_H(g)$ and $a_F(g)$, and wages are w_H and w_F. The standard DFS ordering is summarized by

$$A(g) := \ln \frac{a_F(g)}{a_H(g)}, \qquad A'(g) < 0. \tag{2.1}$$

The function $A(g)$ is Home's relative productivity schedule. Low-g goods are those in which Home has stronger production comparative advantage. If $\varpi := \ln(w_H/w_F)$ denotes the log relative wage, then a frictionless DFS model would simply compare $A(g)$ to ϖ and obtain one production cutoff. In the present note, the wage ratio and the import volume of each traded good are taken as primitives in the main text, so the model should be read as a *conditional* DFS allocation problem whose purpose is to isolate the transport mechanism cleanly. A full demand-and-trade-balance closure can be added later without changing the routing logic.

Equilibrium. Consider now a good g that is imported into Home. To deliver shipment volume m_g, the shipper allocates flow over route locations $x \in [0,1]$, where the midpoint $x = \frac{1}{2}$ is the best corridor. Moving away from the midpoint means using a less direct or less efficient route. Let $f_g(x) \geq 0$ be the density of shipments assigned to route x, so that $\int_0^1 f_g(x)\,\mathrm{d}x = m_g$. A unit shipped on route x incurs intrinsic cost $\tau_0 + \beta_g(x - \frac{1}{2})^2$, where τ_0 is the baseline transport wedge and $\beta_g > 0$ measures how quickly route quality deteriorates away from the midpoint for good g. Congestion is quadratic in local flow and inversely related to the route-capacity profile $K(x) > 0$. Writing $\gamma > 0$ for the congestion parameter, the transport planner solves

$$m_g T_g = \min_{f_g(\cdot)\geq 0} \left\{ \int_0^1 \left[\tau_0 + \beta_g\big(x - \tfrac{1}{2}\big)^2\right] f_g(x)\,\mathrm{d}x + \frac{\gamma}{2} \int_0^1 \frac{f_g(x)^2}{K(x)}\,\mathrm{d}x \right\}, \qquad \int_0^1 f_g(x)\,\mathrm{d}x = m_g. \tag{2.2}$$

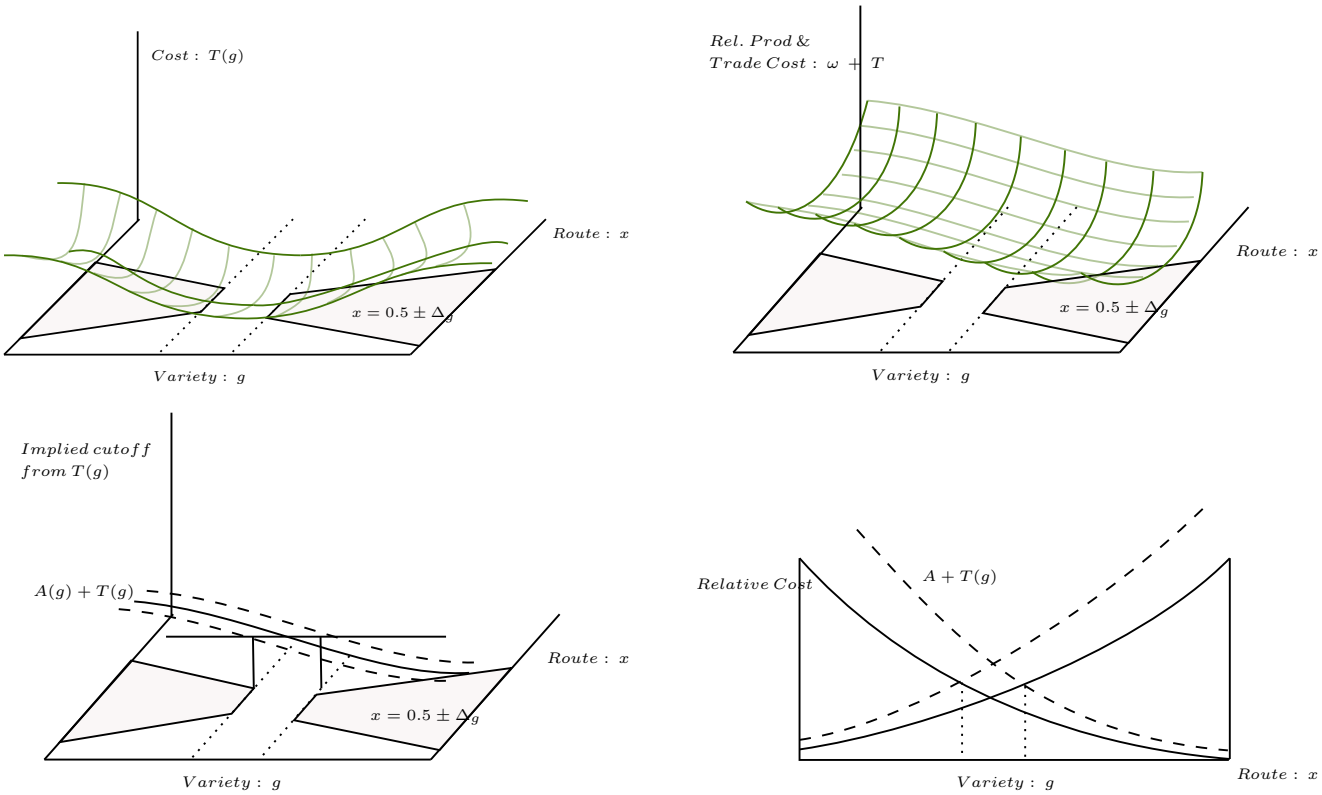


Figure 3. *From route fields to the DFS cutoff.* The upper-left panel shows the transport-cost field generated by the route-allocation problem. For each imported good, the midpoint route is the least-cost corridor and congestion spreads traffic over the active interval $x \in [\frac{1}{2} - \Delta_g, \frac{1}{2} + \Delta_g]$. The upper-right panel adds the production side and shows Home's delivered relative-cost field once the endogenous transport wedge is layered on top of relative productivity. The lower-left panel shows how that route-specific transport wedge projects back into the goods dimension and determines the implied cutoff region. The lower-right panel is the reduced-form DFS representation after route choice has been integrated out. In the benchmark, the outer trade allocation is summarized by the delivered-cost schedule $A(g) + T(g)$ rather than the production schedule $A(g)$ alone.

The minimized object T_g is the per-unit transport cost of imported good g. Under the maintained symmetry assumptions in Appendix A.3, the support of the optimal route density is an interval around the midpoint. Intuitively, congestion prevents all shipments from piling onto the single best corridor, so each imported good ends up using an *area* of routes centered on $x = \frac{1}{2}$. The figure organizes the mechanism in four steps. The upper-left panel shows the transport-cost field that determines the active route band for each good. The upper-right panel overlays that transport wedge on Home's relative production schedule and shows the resulting delivered relative-cost field. The lower-left panel highlights how the route-specific transport wedge translates into the cutoff region over goods. The lower-right panel then collapses the routing problem back into the standard DFS representation, where the relevant one-dimensional comparison is built from the delivered-cost schedule rather than the production schedule alone.

In the analytically tractable benchmark, route capacity is constant, $K(x) \equiv K_0$. The route problem

then collapses to two closed-form objects:

$$\Delta_g = \left(\frac{3\gamma m_g}{4K_0\beta_g}\right)^{1/3}, \qquad T_g = \tau_0 + C_0\left(\frac{\gamma}{K_0}\right)^{2/3}\beta_g^{1/3}m_g^{2/3}, \qquad C_0 := \frac{3}{5}\left(\frac{3}{4}\right)^{2/3}. \tag{2.3}$$

The interpretation is straightforward. High-volume goods use a wider set of routes because congestion makes concentration more expensive. High-capacity networks keep traffic closer to the core corridor and lower transport costs. Goods with larger β_g are more sensitive to route quality, so they pay higher transport costs but use a narrower route interval. The closed-form formula also yields a sharp elasticity: $-\partial T_g/\partial \ln K_0 = \frac{2}{3}(T_g - \tau_0)$. A one-percent increase in capacity lowers the variable component of the transport wedge by exactly two-thirds of its current value. This turns out to be the key local margin behind the policy results. Detailed derivations of the route-support condition, the constant-capacity formulas, and the associated comparative statics are collected in Appendices A.3 and A.4.

Once the routing block is embedded in DFS, the outer trade model remains conceptually unchanged. The only difference is that delivered comparative advantage now depends on production efficiency *and* the transport wedge. In the Home market, Home compares its domestic unit cost $w_H a_H(g)$ to the imported cost $w_F a_F(g)e^{T_{FH}(g)}$. In the Foreign market, the same comparison runs in reverse, with Home paying the export wedge $T_{HF}(g)$. The resulting Home export cutoff $\underline{g}$ and import cutoff $\overline{g}$ are therefore defined by

$$A(\underline{g}) - T_{HF}(\underline{g}) = \varpi, \qquad A(\overline{g}) + T_{FH}(\overline{g}) = \varpi. \tag{2.4}$$

Goods below $\underline{g}$ are exported by Home; goods above $\overline{g}$ are imported by Home; goods in the middle are nontraded. The transport block therefore widens the familiar nontraded band around the frictionless DFS cutoff. When route costs rise, the import cutoff moves upward and the export cutoff moves downward, so trade contracts from both sides. When capacity rises, the reverse happens: the nontraded band narrows and trade expands. The exact cutoff equations and their comparative statics are developed in Appendix A.5.

Equilibrium with Risk. Up to this point the transport problem is deterministic. Risk enters through the availability of route capacity. The model treats risk as uncertainty about usable transport technology rather than uncertainty about demand or post-routing prices. Baseline capacity is installed ex ante, but in a bad state only a fraction of that installed capacity remains usable. This is a reduced-form way to capture the kinds of disruptions that motivate the paper: port closures, sanctions, outages, blockages, labor disruptions, or broader geopolitical events that make parts of the network unavailable precisely when imported goods need them. The timing is straightforward. Baseline capacity is chosen before the state is known. The realization of route availability is then observed, and shippers reroute ex post taking the surviving capacity as given. Foreign dependence matters because imported goods rely on different route locations with different intensities, so a loss of usable capacity is not a uniform shock to trade. Its welfare effect depends on which goods rely on the affected part of the network. Formally, baseline capacity is K_0, but in state ω only a fraction $\xi(\omega) \in (0, 1]$ is available, so effective capacity is $K_0\xi(\omega)$. In each state, the routing problem is identical to the deterministic benchmark with K_0 replaced by $K_0\xi(\omega)$,

so the state-contingent route width and transport wedge are

$$\Delta_g(\omega) = \left(\frac{3\gamma m_g}{4K_0\xi(\omega)\beta_g}\right)^{1/3}, \qquad T_g(\omega) = \tau_0 + C_0\left(\frac{\gamma}{K_0\xi(\omega)}\right)^{2/3}\beta_g^{1/3}m_g^{2/3}. \tag{2.5}$$

Bad states have two simultaneous effects. They raise delivered trade costs and they widen the route area that must be used to sustain a given volume. Expected transport cost therefore exceeds the transport cost implied by average capacity alone. In the common-shock benchmark, expected cost depends on the single moment $\mathbb{E}[\xi^{-2/3}]$, so the reliability premium has a particularly clean interpretation: it is the convexity penalty generated by scarce-capacity states. A small-risk approximation, reported in Appendix A.6, makes this especially transparent by showing that the premium depends on the coefficient of variation of route availability.

Policy. Finally, we turn towards the optimal policy to mitigate risk arising from the transportation infrastructure. Policy is chosen before route availability is realized. The government therefore faces an ex-ante problem. It chooses a subsidy or capacity profile knowing that routing and sourcing will adjust only after the state of the network becomes known. Capacity policy serves two roles at once. In normal states it reduces congestion and lowers the transport wedge. In bad states it limits the rise in transport costs by making the network less vulnerable to capacity shortfalls. A convenient formulation is

$$\max_{s(\cdot)} \; (1-p^c)V_H^0\big(s\big) + p^c\,\mathbb{E}\big[V_H^1(s,\omega)\big] - \int_0^1 \mathcal{C}_x\big(s(x)\big)\,\mathrm{d}x, \tag{2.6}$$

where V_H^0 denotes welfare in the normal state, V_H^1 denotes welfare once route availability has been realized and shippers have rerouted ex post, p^c is the probability of the disruptive state, and $\mathcal{C}_x(s(x))$ is the fiscal cost of subsidizing capacity at route location x. The broad first-order condition is easiest to interpret by separating aggregate from local objects. Let $\mathcal{T}_H$ denote the aggregate import-transport index for Home, defined locally as the expenditure-weighted average of the transport wedges on the imported range. Let $\eta_H^{V,T}$ denote the macro elasticity of expected indirect utility with respect to this aggregate transport index. Under log-separable preferences, this elasticity is simply the expected imported expenditure share, which we denote by α_M. At the local level, let $\Lambda_H(x)$ be the reliance of imported expenditure on route location x, let $\mathcal{B}_H(x)$ be the risk-adjusted bottleneck severity at that location, and let $\chi_K(x)$ be the pass-through from the subsidy into usable capacity. With that notation, the optimal subsidy rule can be written as

$$\mathcal{C}_x'(s^\star(x)) = p^c\,\eta_H^{V,T}\,\Lambda_H(x)\,\mathcal{B}_H(x)\,\chi_K(x). \tag{2.7}$$

This is the broad policy statement of the model. A lane deserves more support when crises are more likely, when aggregate welfare is more sensitive to imported transport costs, when a larger share of imports relies on that lane, when the lane is a more severe risk-adjusted bottleneck, and when the subsidy has stronger pass-through into usable capacity.

In the analytically tractable benchmark with a uniform capacity subsidy, the general policy rule

becomes especially sharp. Because the deterministic transport wedge is homogeneous of degree $-2/3$ in capacity, the exact marginal value of uniform capacity can be written directly in terms of the variable component of the transport wedge. The exact scalar policy condition is

$$\mathcal{C}'(s^\star) = p^c \, \chi_K \, \mathbb{E}\left[\int_{\overline{g}(\omega)}^{1} b(g) \, \frac{2}{3}\big(T_g(\omega) - \tau_0\big) \, \mathrm{d}g\right]. \tag{2.8}$$

This formula is already quite informative: the marginal value of capacity is proportional to the variable part of the transport wedge on the imported range. For exposition, it is also useful to factor this rule into a macro welfare term, a reliability premium, and an aggregate exposure term. Under a fixed-import-set or small-risk approximation one obtains

$$\mathcal{C}'(s^\star) \approx p^c \, \alpha_M \, \Pi \, \Omega_H \, \chi_K, \qquad \Pi := \mathbb{E}[\xi^{-2/3}], \tag{2.9}$$

where Ω_H is the transport-intensity term that depends on the imported goods' β_g, shipment volumes m_g, and baseline congestion. The exact formula in (2.8) is the correct benchmark statement, while (2.9) is the most compact factorized approximation. The derivation of the welfare envelope, the aggregate elasticity objects, and the policy decomposition is given in Appendices A.7 and A.8.

In what follows we present detailed micro data on shipping that, in conjunction with a quantitative multi-country extension of this model, can speak to the elasticities that pin down the optimal policy. The appendix contains the detailed derivations underlying the benchmark formulas used here.

3 Global Shipping Fleet and Foreign Exposure

This section documents the structure of the global containership fleet and how its concentration shapes foreign exposure across countries. We begin by characterizing the global distribution of containership capacity across builder countries, operator nationalities, ship sizes, and vessel vintages. We then relate these global patterns to the fleets that serve U.S. trade in order to measure how global shipbuilding concentration translates into country-level exposure.

Our focus is on containerships because they carry a large share of world trade by value and represent the core mode of long-distance maritime transport. The containership segment is well suited for empirical analysis: its global registry is comprehensive, vessel characteristics are systematically recorded, and ships operate in integrated networks that span multiple regions and trade lanes. These features allow us to trace how global shipbuilding and fleet ownership map into the composition of vessels serving individual trade flows.

The evidence proceeds in two steps. First, we document key facts about the global fleet, including its degree of geographic concentration and the distribution of capacity across ship types. Second, we construct a linked dataset that matches vessel-level records to U.S. import shipments, which permits a detailed measurement of the builder and operator origins of the ships that carry U.S. trade. This approach allows us to quantify how global fleet concentration translates into exposure across routes, ship sizes, and cohorts.

3.1 Data

Our empirical analysis uses two complementary sources that together provide vessel-level coverage of the global containership fleet and shipment-level coverage of U.S. imports.

Clarksons Shipping Intelligence Network (SIN). Clarksons Shipping Intelligence Network (SIN) is a vessel-level registry that covers the global containership fleet. Each vessel is uniquely identified by its International Maritime Organization (IMO) number and includes detailed information on builder country, operator nationality, ship size, vessel vintage, and recorded port rotations. The port-rotation field lists the ports that a vessel regularly serves, allowing us to identify the countries it provides service to on an ongoing basis. For each country or region, we compute the total TEU capacity of all vessels that appear in its port rotations, assigning to each country the full capacity of every vessel that regularly calls at its ports. To maintain consistency, we restrict the sample to active containerships in regular service and exclude vessels listed as idle, laid up, or under repair. Our analysis focuses on the largest segment of the fleet, defined as the roughly 4,000 active containerships with capacities of about 2,000 TEU or more. The sample covers the majority of global TEU capacity and includes all vessels that operate on mainline international services.

Panjiva U.S. Import Data. Panjiva provides shipment-level records for all containerized imports entering U.S. ports. Each record reports the vessel name, arrival port and date, and detailed shipper and consignee information. We match these shipments to Clarksons SIN vessels using IMO numbers and port–date information, which yields a linked dataset that assigns each U.S. import shipment to a uniquely identified vessel with known characteristics. To our knowledge, this is the first dataset that connects the global containership registry to the complete universe of U.S. import transactions.

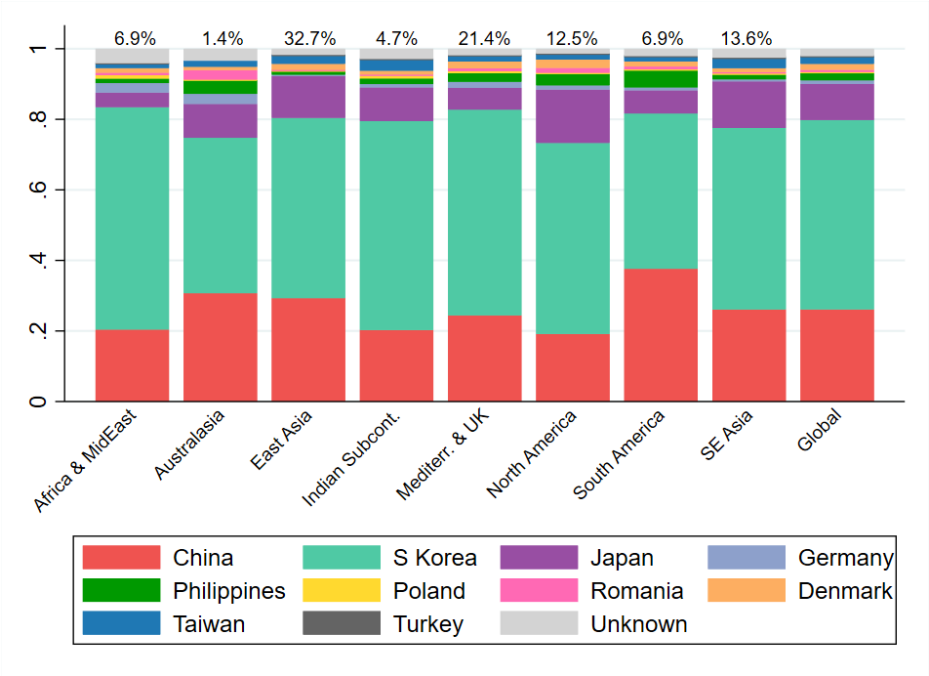
Using the matched data, we compute TEU-weighted measures of U.S. exposure to foreign fleets. For each builder country and operator nationality, we calculate the share of U.S. import capacity carried on vessels of that origin and decompose these measures by ship size and vessel vintage. This construction provides a detailed characterization of the foreign supply base underlying U.S. containerized imports.

3.2 Global Distribution of Ships

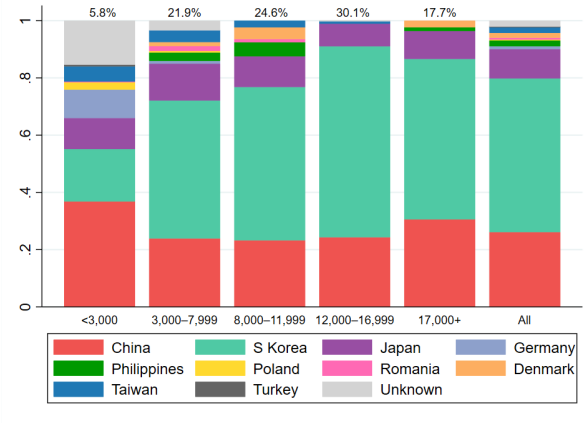
We begin by documenting the global distribution of containership capacity across builder countries. Figure 4 summarizes these patterns using vessel-level data from Clarksons SIN for the active fleet in 2025. Figure 4a reports the overall distribution of capacity by builder country, Figure 4b examines how this distribution varies by ship size, and Figure 4c examines variation by vessel vintage. Together, these panels provide a picture of the geographic concentration of shipbuilding and how it differs across segments of the fleet.

Figure 4a shows that global containership capacity is highly concentrated in a small number of builder countries. South Korea accounts for approximately 54 percent of total TEU capacity and China for an additional 26 percent, implying that 80 percent of the world's active fleet was built in these two countries. With Japan, these East Asian countries account for more than 90 percent of the global fleet.

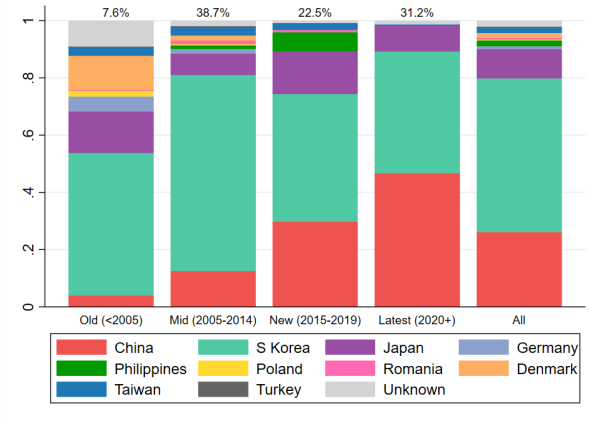
Figure 4. Global Distribution of Containerships by Builder Country, Size, and Age



(a) By Regions



(b) By Vessel Sizes



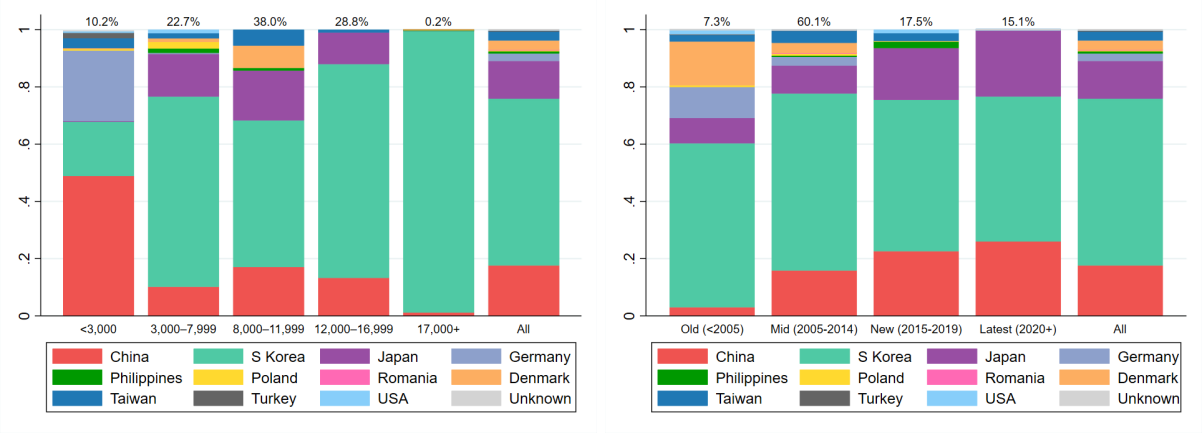
(c) By Vessel Vintages (Age)

Notes: Each bar represents the share of global containership capacity (in TEUs) built in each region, based on vessel-level data from Clarksons SIN (2025 cross-section). Figure 4a reports global shares, Figure 4b disaggregates by ship size, and Figure 4c by vintage. The rightmost column in each figure represents the overall distribution ("Global" and "All"). South Korea and China jointly account for more than 80 percent of total capacity, with dominance strongest for large and new vessels.

Germany, Romania, and other builders contribute much smaller shares, reflecting the dominance of East Asian shipyards in the modern containership market.

Figure 4b disaggregates the global distribution by ship size. The concentration observed in Figure 4a is strongest among larger vessels: South Korea and China together build more than 80 percent of ships above the mainline size thresholds. Their combined share declines to roughly 60 percent among smaller feeder vessels, where production is more dispersed across a wider set of builder countries. This pattern

Figure 5. U.S. Import Fleet Composition by Builder Country, Size, and Vintage



(a) By Size **(b)** By Age

Notes: Each bar shows the TEU-weighted share of U.S. import volumes carried by ships built in each country, by vessel size (Figure 5a) and vintage (Figure 5b). The rightmost column in each figure represents the overall distribution ("Global" and "All"). Data combine Clarksons SIN (vessel registry, 2025) with Panjiva U.S. import records matched by IMO identifiers and port–date information.

reflects the scale-intensive nature of constructing large containerships and the competitive advantage of the largest East Asian yards in these segments.

Figure 4c presents the distribution by vessel vintage. For ships delivered since 2020, China and South Korea together account for nearly 90 percent of global capacity additions, with China responsible for about 47 percent and South Korea for about 43 percent. In contrast, vessels built before 2005 display a more balanced distribution across builder countries, indicating that China's rise in shipbuilding is both recent and rapid. The vintage patterns therefore highlight how the current composition of the global fleet reflects a shift in shipbuilding toward East Asia over the past two decades.

These patterns establish that the global containership fleet is supplied primarily by a small set of East Asian shipbuilding countries, with stronger concentration among larger and newer vessels. Because ships built in these countries appear on major trade routes worldwide, the global distribution of shipbuilding plays a central role in shaping the fleets that serve international trade, including U.S. imports. We summarize these patterns in the following stylized fact:

Stylized Fact 1. *80% of the global containership fleet is built in China and South Korea, with their dominance strongest for large and new ships, while smaller vessels remain more dispersed across other builder countries.countries.*

3.3 U.S. Reliance on Foreign Ships

We next examine the extent to which U.S. containerized imports rely on foreign vessels. Using the matched Clarksons–Panjiva dataset, we compute the TEU-weighted composition of the ships that carry U.S. trade, distinguishing builder countries, operator nationalities, ship sizes, and vessel vintages. Figures 5 and 8 in Section 3.5 summarize these patterns.

Figure 5a reports the builder-country distribution of vessels serving U.S. imports by ship size. The vast majority of U.S. import capacity is carried on vessels built abroad, with East Asian yards accounting for most of the ships in every size category. South Korea and Japan dominate the largest size classes, reflecting their leading role in constructing large mainline vessels, while China has a more prominent presence among feeder and mid-sized ships.

Figure 5b presents the distribution by vessel vintage. Among ships delivered since 2020, nearly half of the vessels serving U.S. trade were built in China, compared with a very small share among ships built before 2005. This pattern reflects the broader global shift in shipbuilding toward Chinese yards over the past decade and indicates that China has become an increasingly important source of the capacity that will shape U.S. trade in coming years.

These patterns motivate the following stylized fact, which highlights the global integration of Chinese-built vessels into U.S. import logistics:

Stylized Fact 2. *U.S. imports are overwhelmingly carried on foreign-built ships, with East Asian yards (South Korea, China, Japan) dominating across both size and vintage. China has rapidly gained ground in smaller and newer vessels, while Korea and Japan remain central in larger and mid-aged ships.*

3.4 Ship Prices by Size and Origin

We next examine the cost side of global shipbuilding and assess whether the concentration documented above is associated with systematic differences in vessel prices. Using Clarksons SIN data on ship orders and contract prices, we study how the cost per unit of capacity, measured in dollars per TEU, varies with vessel size across builder countries. The sample consists of all containerships in the 2025 global orderbook for which both price and capacity information are available. Figure 6 summarizes these patterns.

Figure 6 plots price per TEU against vessel size. Vessel prices decline sharply as capacity increases, indicating pronounced economies of scale in containership production. At comparable sizes, Chinese-built ships tend to be priced below vessels built in South Korea, suggesting that China's recent expansion in shipbuilding has been accompanied by cost advantages at given capacity levels.

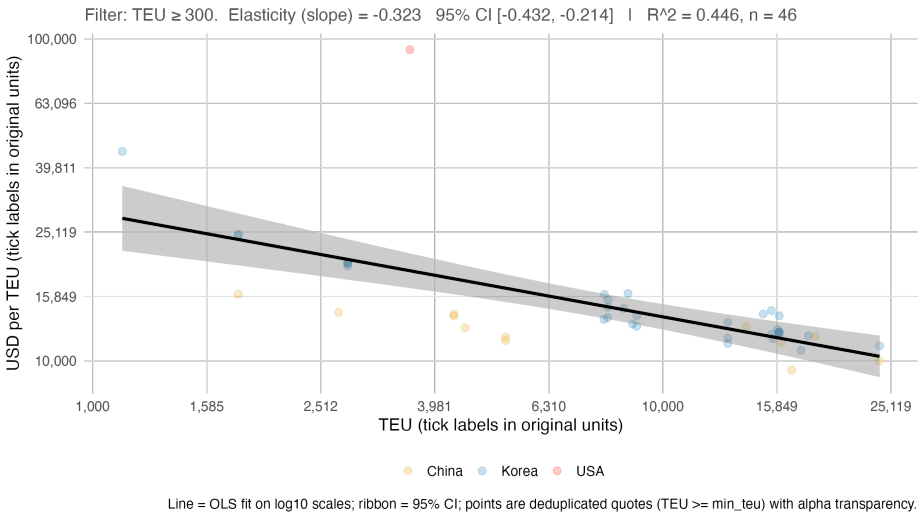
Taken together, these patterns show that the global concentration of shipbuilding is associated with meaningful cost differences across producers. The presence of strong scale economies and lower contract prices for Chinese-built vessels highlights the efficiency gains available from access to large foreign shipyards, complementing the earlier evidence on the dominant role of East Asian builders in supplying the modern containership fleet. We summarize these patterns in the following stylized fact:

Stylized Fact 3. *Both China and South Korea build highly efficient ships, with prices per TEU falling sharply with vessel size. Chinese-built ships are generally cheaper.*

3.5 Foreign-Operated Ships

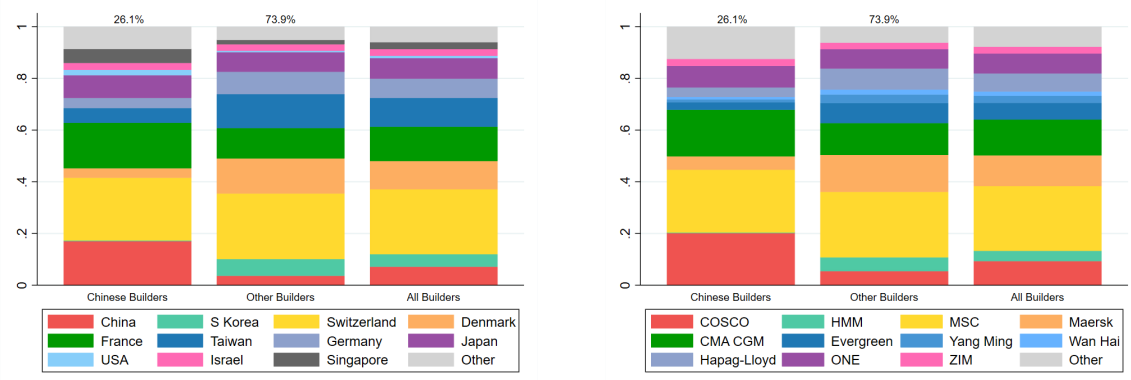
The patterns we have established so far describe the origins of the vessels that carry global trade and U.S. imports, but builder nationality does not necessarily correspond to operating control. Shipping

Figure 6. Ship Prices per TEU, 2025 Containership Orderbook



Notes: Each point represents a containership on order in 2025 with reported capacity and contract price, drawn from Clarksons SIN. The fitted relationship shows that unit prices decline with vessel size, reflecting strong scale economies in shipbuilding. Chinese-built ships are generally cheaper at comparable capacities.

Figure 7. Global Operator Nationality of Ships



(a) Operator Country Distribution **(b)** Operator Distribution

Notes: Figure 7a shows the distribution of countries of container shipping operators globally while Figure 7b shows distribution of the container shipping operators. In both figures, the left panel shows the global distribution of operator nationalities among Chinese-built vessels; the middle panel shows the distribution among vessels not built by China; and the right panel shows the overall operator distribution. Chinese-built ships are predominantly operated by non-Chinese carriers headquartered in Europe or other East Asian countries. Data from Clarksons SIN.

services are provided by carriers that lease, own, or deploy vessels in global networks that span multiple regions, and the firms that operate these ships often differ from the firms or countries that built them. To distinguish these two dimensions of foreign exposure, Figures 7 and 8 report the distribution of vessel operator nationalities and operator firms at both the global and U.S. levels respectively.

The left column of Figure 7a shows that a vast majority of Chinese-built ships (78%) are operated by countries other than China, including Switzerland (24.8%) and France (18.3%). Chinese operators account for about a fifth of Chinese-built ships (22%). In contrast, South Korea plays a much smaller

role as a vessel operator (4.6% of overall volumes, rightmost column of Figure 7a), especially relative to China, which is both a major builder and a significant operator of the global fleet. In terms of the specific firms, major global container operators MSC (Switzerland), CMA CGM (France), and ONE (Japan) account for 51.7% of Chinese-built ships (left column, Figure 7b). Overall, operator countries and firms are far less concentrated than builder countries. Because operator firms are highly visible in global shipping, particularly through their branded container fleets, it is important to highlight both sides of the market: while shipbuilding is dominated by a small number of countries, vessel operation is considerably more dispersed across nations and firms.

Figure 8 reports the distribution of operator nationalities for the vessels serving U.S. trade. The leftmost column in Figure 8a reports operator nationalities for the subset of vessels serving U.S. imports that were built in China. Despite their builder origin, these vessels are predominantly operated by carriers headquartered outside China, including Switzerland, France, Denmark, and the United States. Notably, U.S. operators make up 17.4% of Chinese-built ship volumes. Chinese operators, similar to the global distribution in Figure 7, represent less than a third of activity within this group (about 31.3%). This pattern indicates that the presence of Chinese-built vessels in U.S. trade largely reflects China's role in global shipbuilding rather than a concentration of operating control. It also highlights that U.S. policies targeting Chinese-built ships risk unintended consequences, as a substantial share of these vessels are operated by countries other than China, even including U.S. operators.

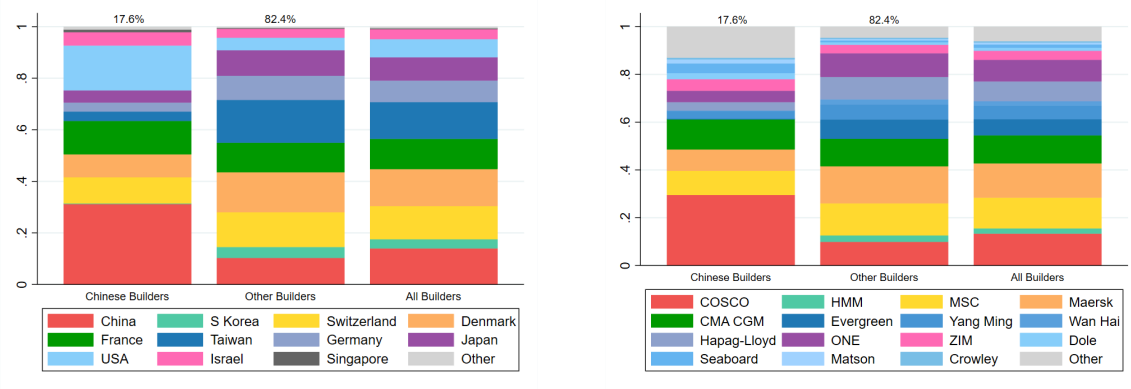
Figure 8b presents operator nationalities for all vessels serving U.S. imports, regardless of builder origin. The distribution is broadly similar to that in Figure 8a: major global carriers such as Maersk, MSC, CMA CGM, Hapag-Lloyd, and ONE dominate U.S.-serving routes. The small share accounted for by Chinese operators reinforces the observation that shipping operations remain globally diversified even as shipbuilding output is highly concentrated. The contrast between the builder and operator distributions establishes that ship construction and fleet operation are distinct activities within the container shipping sector.

Taken together, these patterns show that U.S. containerized imports are carried almost entirely on foreign fleets, with a high degree of concentration in East Asian shipbuilding and a globally diversified set of operators. The U.S. reliance on foreign-built vessels spans ship sizes and vintages, and the operator patterns underscore the integrated nature of the global shipping industry. We summarize this in the following stylized fact:

Stylized Fact 4. *Globally, the majority (78%) of Chinese-built vessels are operated by non-Chinese firms and countries. For Chinese-built vessels serving U.S. imports, this share remains high at 68.7% with a sizeable portion of operated by U.S. carriers. These patterns show that while the shipbuilding industry is highly concentrated, the operational control of these ships is far more dispersed. Chinese-built ships are deeply integrated into international shipping networks, with their usage extending well beyond Chinese operators.*

Thus, the evidence documented in this section characterizes a global shipping sector that is geographically concentrated, globally interconnected, and marked by substantial economies of scale. South Korea and China supply the majority of the world's containership capacity, and the vessels that carry U.S. imports are almost entirely foreign built and foreign operated, with large East Asian yards dominating

Figure 8. Operator Nationality of Ships Serving U.S. Imports



(a) Operator Country Distribution

(b) Operator Distribution

Notes: Figure 8a shows the distribution of countries of container shipping operators servicing U.S. imports while Figure 8b shows distribution of the container shipping operators. In both figures, the left panel shows the distribution of operator nationalities among Chinese-built vessels serving U.S. trade; the right panel shows the overall operator distribution. Chinese-built ships are predominantly operated by non-Chinese carriers headquartered in Europe or other East Asian countries. Data combine Clarksons SIN (vessel registry, 2025) with Panjiva U.S. import records matched by IMO identifiers and port–date information.

both current and forthcoming fleet cohorts. At the same time, the cost patterns in shipbuilding indicate sizable efficiency gains associated with access to these producers. These facts motivate the central question for the next section: how does reliance on foreign-controlled shipping capacity shape the prices and routes that govern U.S. trade, and how do these efficiency gains compare to the exposure risks that arise in such a concentrated global supply base?

4 Quantitative Model of Trade with Shipping and Risk

We develop a model of international trade in which goods move across a global shipping network, and in which shipping services are supplied by fleets owned in different countries and subject to risk. Fleets make capacity decisions ex-ante, before uncertainty about their availability on different trade links is realized, and shipping services are allocated ex-post through operators that rent available capacity and supply transport on a directed network. Shipping costs are determined endogenously from the interaction of routing, congestion, and the mix of fleets serving each route in each state of the world. This structure allows the model to capture how exposure to foreign fleets and to shocks in their availability affects delivered trade costs, trade flows, and welfare. We now describe the environment, the behavior of households, firms, operators, and fleet producers, and the definition of competitive equilibrium.

4.1 Countries and Locations

The world consists of M countries indexed by $m = 1, \ldots, M$, each containing a finite set of production and consumption locations (nodes) denoted by $\mathcal{N}_m$. Let $\mathcal{N} = \bigcup_{m=1}^{M} \mathcal{N}_m$ denote the set of all nodes. Locations are connected by a directed transport network $G = (\mathcal{N}, \mathcal{E})$, where each directed edge

$e = (i \to j) \in \mathcal{E}$ represents a feasible shipping link from origin i to destination j. The network need not be symmetric, and the presence of $(i \to j)$ does not imply the availability of $(j \to i)$. The geographic structure determines which pairs of nodes can be directly connected by shipping services and underpins the routing of goods throughout the network.

4.2 Households

Each country m is populated by a representative household with total labor endowment $\bar{L}_m$. Labor is perfectly mobile across nodes within each country but immobile across countries. Hence any feasible allocation must satisfy

$$\sum_{i \in \mathcal{N}_m} L_i = \bar{L}_m,$$

and free mobility within a country implies that real income is equalized across all nodes $i \in \mathcal{N}_m$.

At location j, households consume a CES composite of the varieties produced across all nodes,

$$C_j = \left(\sum_{i \in \mathcal{N}} \phi_{ij}^{1/\sigma} q_{ij}^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, \qquad \sigma > 1,$$

where ϕ_{ij} are taste weights. Location j provides an amenity level u_j, which may depend on local population. With local wage w_j and CES price index P_j, indirect utility is $U_j = u_j w_j / P_j$, and mobility within the country requires $U_j = W_m$ for all $j \in \mathcal{N}_m$, where W_m is the country-wide welfare level.

4.3 Production

Each location $i \in \mathcal{N}$ produces a unique variety using local labor. Output is given by

$$Y_i = A_i L_i,$$

where A_i denotes labor productivity. Productivity may depend on local population:

$$A_i = \bar{A}_i L_i^{\alpha_{m(i)}},$$

where $\bar{A}_i$ is an exogenous productivity level and $\alpha_{m(i)}$ governs the strength of agglomeration or congestion effects within country $m(i)$. Under perfect competition, firms take wages as given and price their variety at marginal cost, which implies a factory-gate price

$$p_i = \frac{w_i}{A_i}.$$

Goods produced at i may be consumed locally or shipped to other locations, subject to the shipping technology described below.

4.4 Shipping

Shipping services move goods across the directed transport network and are produced by two sets of agents: shipping fleet operators, who assemble services on individual links, and shipping fleet producers, who supply the underlying capacity. Shipping costs are therefore endogenous and reflect both the mix of fleets serving each route and the shadow cost of congestion on the network.

Transport links and services. Each directed edge $e = (i \to j)$ in the network represents a shipping link on which transport services may be supplied. Let $S_e(\Omega)$ denote the quantity of services provided on edge e in state Ω. Generalized shipping costs on an edge combine a technological component and an endogenous user-price component, both of which will be defined below.

Risk and timing. The availability of shipping fleets on different edges is subject to uncertainty. At the beginning of the period, shipping fleet producers choose capacity levels $\{f_n\}_{n=1}^{N_f}$ before uncertainty is realized. A state Ω is then drawn, where each indicator $\Omega_{e,n} \in \{0, 1\}$ specifies whether fleet origin n is able to supply services on edge e. Once Ω is known, operators rent available capacity at rental rates $\rho_n(\Omega)$, shipping services are supplied, and goods are routed across the network. Because capacity is chosen ex-ante and remunerated ex-post, rental income is state-contingent and embeds risk for producers.

Shipping fleet operators. For each edge $e = (i \to j)$, a competitive operator rents capacity from fleet origins $n = 1, \ldots, N_f$ in state Ω. Let $s_{e,n}(\Omega)$ denote the quantity of capacity from origin n rented on edge e. Capacity is combined using a CES technology with elasticity of substitution $\eta > 1$:

$$S_e(\Omega) = \left(\sum_{n=1}^{N_f} \Omega_{e,n}\, \omega_{e,n}^{1/\eta}\, s_{e,n}(\Omega)^{(\eta-1)/\eta} \right)^{\eta/(\eta-1)}.$$

Supplying $S_e(\Omega)$ units of service generates a congestion cost

$$\Psi_e(S_e(\Omega)) = \frac{\zeta_e}{1+\delta_e}\, S_e(\Omega)^{1+\delta_e}, \qquad \zeta_e \geq 0,\ \delta_e \geq 0.$$

Taking rental rates $\rho_n(\Omega)$ and output price $p_e(\Omega)$ as given, the operator chooses $\{s_{e,n}(\Omega)\}$ to maximize profits:

$$\max_{\{s_{e,n}(\Omega)\}}\ p_e(\Omega)\, S_e(\Omega) - \sum_{n=1}^{N_f} \rho_n(\Omega)\, s_{e,n}(\Omega) - \Psi_e(S_e(\Omega)).$$

Under perfect competition, equilibrium profits are zero. The implied per-unit user price is

$$p_e(\Omega) = \left(\sum_{n=1}^{N_f} \Omega_{e,n}\, \omega_{e,n}\, \rho_n(\Omega)^{1-\eta} \right)^{1/(1-\eta)} + \zeta_e\, S_e(\Omega)^{\delta_e}.$$

Shipping fleet producers. Shipping fleet producers supply the capacity that operators rent on each shipping link. Producer n chooses a capacity level f_n ex-ante, before uncertainty about route availability is realized. Supplying f_n units of capacity requires the convex investment cost

$$\frac{1}{2}\,\mu_n f_n^2, \qquad \mu_n > 0.$$

After the state of the world Ω is realized, the portion of fleet n that is available to operate on edge e is given by the indicator $\Omega_{e,n} \in \{0, 1\}$, and available capacity earns the rental rate $\rho_n(\Omega)$. Shipping fleet producer n, located in country $m(n)$, chooses capacity to maximize expected discounted profits:

$$\max_{f_n \geq 0} \; \mathbb{E}_\Omega \left[\Lambda_{m(n)}(\Omega) \left(\rho_n(\Omega)\, f_n - \tfrac{1}{2}\,\mu_n f_n^2 \right) \right],$$

where $\Lambda_{m(n)}(\Omega)$ is the stochastic discount factor of the fleet-owning country. Capacity is chosen ex-ante but remunerated ex-post, so investment responds to both expected rental revenues and the risk embedded in $\rho_n(\Omega)$.

Generalized shipping costs. Let $\bar{\kappa}_e > 0$ denote the technological component of the shipping cost on edge e and let v_e denote the average value of shipments transported along e. We allow this technological component to be subject to risk: for each edge,

$$\ln \bar{\kappa}_e = \mu_e + \varepsilon_e, \qquad \varepsilon_e \sim \mathcal{N}(0, \sigma_e^2),$$

independently across edges. Given a realization of $\bar{\kappa}_e$, the ad-valorem shipping cost is $p_e(\Omega)/v_e$, and the corresponding generalized iceberg factor is

$$\kappa_e(\Omega) = \bar{\kappa}_e \left(1 + \frac{p_e(\Omega)}{v_e} \right).$$

4.5 Routing and Trade Costs

Shipments can reach a destination through many possible sequences of edges in the transport network. To capture heterogeneous routing preferences across otherwise identical shippers, we assume that each feasible route is subject to an idiosyncratic taste shock. These taste shocks represent variation in convenience and familiarity across routes and induce a probabilistic allocation of traffic across paths. This environment yields the familiar soft minimum aggregation of route costs used in models of random utility.

Routing under uncertainty. A shipper sending a unit from i to j chooses a route before the realization of both the technological shocks $\{\varepsilon_e\}$ and the taste shocks $\{\nu_r\}$. For each feasible route $r \in \mathcal{R}(i \to j)$, let ν_r be an idiosyncratic taste shock drawn independently from a Type I extreme value distribution. Let

$\mathcal{E}(r)$ denote the sequence of edges that constitute route r, and define the random log cost of the route by

$$Z_r \equiv \sum_{e \in \mathcal{E}(r)} \ln \kappa_e.$$

The randomness in Z_r reflects the technological shocks that affect the edges along the route. Shippers are risk averse with respect to these technological shocks and evaluate each route using a CARA objective with parameter $\chi > 0$. With CARA preferences applied to log costs, the resulting evaluation of uncertain transport conditions can be expressed through a certainty-equivalent cost, providing a tractable way to represent risk in routing decisions. The routing problem is therefore

$$\min_{r \in \mathcal{R}(i \to j)} \mathbb{E}\left[\exp\{\chi(Z_r - \nu_r)\}\right],$$

where expectations are taken over both the technological cost shocks and the taste shocks. Under the Type I extreme value assumption for ν_r, with scale parameter governed by a constant $\varphi > 1$.

Certainty-equivalent edge costs. When evaluating uncertain technological conditions on each edge, shippers form expectations over the random log costs generated by the shocks $\{\varepsilon_e\}$. If $\ln \kappa_e$ is normally distributed with mean μ_e^κ and variance s_e^2, then under CARA preferences with parameter $\chi > 0$ the certainty equivalent log cost of edge e is

$$\ln \kappa_e^{CE} \;=\; \mu_e^\kappa \;+\; \frac{\chi}{2}\, s_e^2.$$

The corresponding certainty equivalent generalized cost is therefore

$$\kappa_e^{CE} \;=\; \exp\!\left(\mu_e^\kappa \;+\; \frac{\chi}{2}\, s_e^2\right).$$

These certainty equivalent edge costs summarize the effect of technological uncertainty on perceived transport costs when routing decisions are made. For all routing expressions that follow, the relevant effective edge cost is therefore

$$\kappa_e(\Omega) \;\equiv\; \kappa_e^{CE}(\Omega).$$

Recursive trade costs. Integrating over the idiosyncratic taste shocks $\{\nu_r\}$ in the route-selection problem yields the standard logit structure, under which the overall cost of reaching a destination j from an origin i satisfies a soft-minimum recursion. In particular, letting $\kappa_{ik}(\Omega)$ denote the generalized cost of shipping along edge $(i \to k)$ in state Ω, the resulting aggregation over routes implies

$$\tau_{ij}(\Omega)^{\,1-\varphi} = \mathbf{1}\{i = j\} + \sum_{k \in N(i)} \kappa_{ik}(\Omega)^{\,1-\varphi}\, \tau_{kj}(\Omega)^{\,1-\varphi}, \qquad \varphi > 1.$$

This expression reflects the idea that reaching j from i amounts to choosing a first step $i \to k$ and then continuing optimally from k to j. The recursion summarizes the entire distribution of route-specific shocks into a single, tractable representation of the effective bilateral trade barrier between i and j.

Routing probabilities. Given this structure, the probability that a shipment from i to j uses edge $(k \to \ell)$ is

$$\theta_{(k\to\ell)}(i \to j; \Omega) = \frac{\tau_{ik}(\Omega)^{\,1-\varphi}\, \kappa_{k\ell}(\Omega)^{\,1-\varphi}\, \tau_{\ell j}(\Omega)^{\,1-\varphi}}{\tau_{ij}(\Omega)^{\,1-\varphi}}.$$

This expression mirrors multinomial-logit choice probabilities: an edge is more likely to be used if it is cheap relative to alternatives and if it leads to downstream nodes that are themselves well connected to j.

The recursive formulation implies that any change in an edge-level cost $\kappa_{ik}(\Omega)$ affects not only routes that directly use the edge but also all routes for which the edge lies on a path with positive probability. A decline in κ_{ik} reduces the cost of reaching every downstream node in proportion to the accessibility of those nodes from k. Conversely, an increase in costs or a disruption that makes $\Omega_{e,n} = 0$ on certain links raises bilateral iceberg trade costs for all origin–destination pairs that rely on those links. Routing shares $\theta_{(k\to\ell)}(i \to j; \Omega)$ capture the resulting reallocation of traffic across paths. This mapping from edge-level shipping prices to bilateral trade costs is the mechanism through which the shipping network transmits shocks into the spatial trade equilibrium.

4.6 Competitive Equilibrium

We take the price of a numeraire good as given. A competitive equilibrium is a collection of prices and allocations such that, for every state of the world Ω, the following hold:

(i) Households maximize utility and relocate within each country m to equalize real income across all nodes $j \in \mathcal{N}_m$, and the labor constraint $\sum_{i\in\mathcal{N}_m} L_i = \bar{L}_m$ holds for each country.

(ii) Firms at each node i choose labor to maximize profits, produce $Y_i = A_i L_i$, and set $p_i = w_i/A_i$ under perfect competition.

(iii) Shipping fleet operators choose $\{s_{e,n}(\Omega)\}$ on each edge e to maximize profits given prices and congestion; shipping fleet producers choose $\{f_n\}$ ex-ante to maximize expected discounted profits.

(iv) Routing probabilities and bilateral iceberg trade costs satisfy the recursive formulation described in the routing subsection.

(v) Goods markets clear: $Y_i = \sum_j q_{ij}(\Omega)$ for every node i.

(vi) Shipping services markets clear: $S_e(\Omega) = \sum_{i,j} q_{ij}(\Omega)\, \theta_e(i \to j; \Omega)$ for every edge e.

(vii) Fleet capacity markets clear: $\sum_{e\in\mathcal{E}} s_{e,n}(\Omega) = f_n$ for every fleet producer n.

(viii) Labor markets clear: $\sum_{i\in\mathcal{N}_m} L_i = \bar{L}_m$ for every country m.

An equilibrium is any collection of wages, prices, allocations, flows, capacities, and routing probabilities satisfying conditions (i)–$(viii)$.

5 Shipping Shocks in the Global Transport Network

This section develops the theoretical mapping from fleet-specific shocks to shipping prices, route-level iceberg costs, and welfare. The analysis proceeds in three steps. First, we characterize how a change in the rental rate of a given fleet affects the cost of shipping services supplied on each edge of the network. Second, we show how these edge-level cost changes propagate through the routing structure of the transport system and alter delivered trade costs between all origin–destination pairs. Finally, we derive the welfare consequences of small local shipping shocks and discuss the mechanisms that shape their magnitude and distribution. These results provide the analytical foundation for the quantitative exercise in the next section. To focus on the core economic mechanisms, this section considers the deterministic version of the model, abstracting from edge-level cost risk and from uncertainty in fleet availability.

5.1 Fleet-Specific Shocks in the Shipping Sector

Consider any shock that raises the cost of deploying fleet n, represented in the model as a proportional increase in its rental rate ρ_n. Because shipping fleet operators combine inputs from multiple fleets using a CES aggregator, such a change increases the per-unit price of shipping services on every edge e where fleet n contributes to supply. The resulting pass-through into the generalized iceberg cost on edge e is

$$\frac{\partial \ln \kappa_e}{\partial \ln \rho_n} = a_{e,n} \frac{p_e}{p_e + v_e},$$

where $a_{e,n}$ is the cost share of fleet n in the shipping services used on edge e, and $p_e/(p_e + v_e)$ maps the user-price change into an ad valorem shipping-cost change.

For small proportional shocks to the rental rate, this derivative implies the local change

$$\Delta \ln \kappa_e \;\approx\; a_{e,n} \frac{p_e}{p_e + v_e} \Delta \ln \rho_n,$$

which is the mapping used in the quantitative analysis.

This expression shows that a higher fleet-specific rental cost raises route-level shipping costs in proportion to the route's reliance on that fleet and the importance of shipping in shipment values.

5.2 Propagation Through Routing and Network Structure

Changes in edge-level shipping costs propagate through the transport network through both direct and indirect channels. A higher iceberg cost on edge $(i \to k)$ raises the cost of any route that uses that edge with positive probability. Because routing is determined by a discrete-choice structure with route-specific shocks, shippers substitute toward alternative paths in proportion to the relative increase in route costs. As a result, a shock to κ_{ik} affects the bilateral trade costs τ_{ij} for every origin–destination pair whose least-cost or probabilistic routes use $(i \to k)$ in equilibrium.

Formally, the elasticity of the bilateral trade cost with respect to an upstream edge cost is

$$\frac{\partial \ln \tau_{ij}}{\partial \ln \kappa_{ik}} = \theta_{(i\to k)}(i \to j) + \sum_{(\ell\to m)\neq(i\to k)} \theta_{(\ell\to m)}(i \to j)\, \frac{\partial \ln \kappa_{\ell m}}{\partial \ln \kappa_{ik}},$$

where $\theta_{(\ell\to m)}(i \to j)$ denotes the routing probability that a shipment from i to j uses edge $(\ell \to m)$ in equilibrium. The first term captures the direct effect of the shock on routes that include $(i \to k)$, while the second captures the indirect effect that arises because higher costs on one edge shift routing probabilities across all alternative paths. When the network contains close substitutes, these substitution responses mitigate propagation, whereas bottleneck edges amplify it.

These routing elasticities feed directly into the welfare mapping in the next subsection. Welfare depends on delivered prices, and delivered prices aggregate the bilateral trade costs τ_{ij}. The same routing structure that allocates shipments across paths therefore determines how localized shipping shocks translate into aggregate welfare effects.

5.3 Welfare Effects of Transport Cost Changes

Welfare responds to shipping shocks through their impact on delivered trade costs. Combining the routing structure above with the mapping from fleet-specific shocks to edge-level iceberg costs yields a decomposition of the welfare elasticity into a spatial component and a shipping component,

$$\Delta \ln W = \sum_{e'} \left(\frac{\partial \ln W}{\partial \ln \kappa_{e'}} \right) \Delta \ln \kappa_{e'}.$$

The spatial term $\partial \ln W / \partial \ln \kappa_{e'}$ captures how strongly the economy relies on edge e', reflecting its baseline traffic, its position in the network, and the inward and outward access multipliers of its endpoints. The shipping term $\Delta \ln \kappa_{e'}$ summarizes how the policy shock alters edge-level costs.

The shipping response itself operates through three channels. First, there is a direct effect on the taxed or shocked edge, proportional to its reliance on the targeted fleet and its shipping-price share. Second, shocks that alter the rental rates of individual fleets propagate to all edges that use those fleets, generating market-wide spillovers. Third, changes in the cost of individual edges induce rerouting, which can either amplify or dampen the impact on other edges depending on the availability of substitutes and the degree of congestion. These channels together determine $\Delta \ln \kappa_{e'}$ and therefore the aggregate welfare effect.

For small shocks, welfare changes take the following form:

$$\Delta \ln W \approx -\sum_{e'} \omega_{e'}\, \Delta \ln \kappa_{e'},$$

where $\omega_{e'}$ is the expenditure share of trade using edge e', adjusted for routing elasticities. This representation makes clear that welfare losses depend jointly on direct exposure, spillovers through fleet markets, and the network structure governing substitution across routes.

5.4 A Benchmark Case for Welfare Incidence

It is useful to consider a simplified environment that abstracts from spatial externalities in production or amenities and from congestion in the shipping sector. In this case, routing probabilities respond only to relative transport costs, fleet rental rates adjust competitively, and edge-level pass-through reflects exposure to the fleets used on each edge. This benchmark delivers a transparent decomposition of the welfare effects of transport-cost shocks.

Let Ξ_e denote the baseline routing weight on edge e, defined as the probability that a randomly drawn shipment in the economy uses edge e in equilibrium. Let $s_{e|n}$ denote the share of shipping services on edge e supplied by fleet n, and let $\psi_e = p_e/(p_e + v_e)$ denote the shipping-price share on that edge, as in the previous subsection. Finally, let $a_{e,n}$ be the cost share of fleet n in the shipping aggregator on edge e, and let $\phi_{n;e}$ denote the elasticity of the rental rate of fleet n with respect to a shock that affects edge e.

In this benchmark environment, the welfare elasticity with respect to a small change in the shipping cost on edge e associated with fleet n can be written as

$$\frac{d\ln W}{d\ln \kappa_{e,n}} = \underbrace{\Xi_e\, s_{e|n}}_{\text{Direct exposure on the affected edge}} + \underbrace{\sum_{e'} \Xi_{e'}\, \psi_{e'} \sum_{n'} a_{e',n'}\, \phi_{n';\,e}}_{\text{Indirect spillovers through fleet-market linkages}} .$$

The first term captures the direct impact of a higher shipping cost on edge e through the share of shipments that use that edge and the dependence of that edge on fleet n. The second term captures how changes in fleet rental rates propagate the shock to all other edges that use fleets whose costs adjust in response to the original disturbance.

This decomposition highlights the two central channels through which transport-cost shocks affect welfare in the full model. The direct channel reflects the exposure of trade flows to the specific edges and fleets that experience the cost increase. The indirect channel reflects the general-equilibrium interactions operating through fleet markets and routing adjustments.

6 U.S. Shipping Fees on Chinese Ships

We now apply the framework developed above to quantify the welfare effects of the surcharges introduced by the United States on containerships that are Chinese-built or Chinese-operated. This exercise illustrates how fleet-specific cost shocks propagate through the global shipping network and how these shocks translate into aggregate welfare changes. We proceed in four steps. First, we summarize the key institutional features of the policy. Second, we map the statutory fees into the model as shifts in the rental rates of the affected fleets. Third, we combine these shocks with observed lane-level exposure to Chinese-built capacity to compute the resulting route-level shipping-cost changes. Finally, we aggregate these changes to obtain an estimate of the national welfare impact.

6.1 Policy Background

In April 2025, the United States announced two new surcharges under Section 301 of the Trade Act of 1974 that directly target containerships linked to China. These measures were originally scheduled to take effect on October 14, 2025, but their implementation has since been postponed until late 2026.

The first surcharge applies to Chinese-operated ships and is set at $50 per net ton per U.S. port rotation, up to a maximum of five rotations per year. The second applies to Chinese-built ships and is levied per net ton or per container—whichever is higher—with a schedule that was to begin at $18 per net ton in 2025 and rise to $33 by 2028. A vessel that is both Chinese-built and Chinese-operated would be subject to both fees.

These measures extend the scope of trade policy beyond goods and into the maritime logistics sector, reflecting heightened attention to foreign dependence in critical infrastructure. Although justified as national-security provisions, the surcharges function economically as targeted taxes on shipping services. They therefore offer a useful setting for examining how policies affecting specific fleets alter shipping costs, re-route global trade flows, and affect aggregate welfare. The fact that they apply on a per-call basis and scale with ship size makes them especially well suited for the lane-by-lane analysis we undertake below.

6.2 Mapping Surcharges into Model Objects

In the model, a surcharge applied to a particular fleet n enters as an ad valorem tax T_n on the use of that fleet. Because fleet n supplies a differentiated input into shipping services, this tax raises the effective rental rate of the fleet in proportion to its level:

$$\Delta \ln \rho_n = \ln(1 + T_n).$$

A higher rental rate then increases the user price of shipping services on every edge e where fleet n contributes to supply.

The magnitude of this pass-through is governed by the fleet's cost share on the edge and by the importance of shipping charges in total shipment values. Formally,

$$\frac{\partial \ln \kappa_e}{\partial \ln \rho_n} = a_{e,n} \frac{p_e}{p_e + v_e},$$

where $a_{e,n}$ is fleet n's cost share in the CES aggregator of shipping services on edge e, p_e is the per-unit price charged to users of shipping on that edge, and v_e is the average value of shipments transported on that edge. For a finite tax T_n, the resulting change in the iceberg cost on edge e is therefore

$$\Delta \ln \kappa_e \approx a_{e,n} \frac{p_e}{p_e + v_e} T_n.$$

This mapping highlights that the effect of a surcharge depends directly on the reliance of a route on the taxed fleet and on the weight of shipping charges in delivered prices.

To bring the statutory fees into the model, we convert the per-ton and per-container surcharges into

effective ad valorem taxes T_n^{eff} on the corresponding fleets. The conversion uses data on vessel capacities, average loads, and representative shipment values to translate a per-call fee into an average per-container cost increase. Because the statutory fees scale with ship size and with the number of U.S. port calls, the resulting ad valorem rates vary across fleets and over time. Using observed capacity distributions and typical cargo values, we obtain effective tax rates of roughly 48 percent for Chinese-operated ships and between 17 and 32 percent for Chinese-built ships over the 2025–2028 phase-in schedule.

These tax rates summarize the proportional increase in per-container shipping costs faced by importers when shipments are carried on the affected vessels. Combined with the exposure of each shipping lane to the targeted fleets, they provide the key inputs for quantifying the impact of the surcharges on route-level shipping costs and, ultimately, on national welfare.

6.3 Lane-Level Exposure and Welfare Aggregation

The theoretical results in Section 4 imply that the welfare effect of a small change in shipping costs can be expressed as an expenditure-weighted average of route-level price changes, scaled by the share of imports in GDP. If $\Delta \ln \kappa_e$ denotes the proportional change in the iceberg shipping cost on lane e, then

$$\frac{\Delta \text{EV}}{Y} \approx - s_{\text{imp}}^{\text{GDP}} \sum_e w_e \, \Delta \ln \kappa_e,$$

where $s_{\text{imp}}^{\text{GDP}}$ is the share of imports in GDP and w_e is the share of aggregate imports that use lane e. This decomposition mirrors the model's structure: welfare responds to higher shipping costs in proportion to (i) how much those costs raise delivered import prices and (ii) how important imports are in national expenditure.

To implement this expression empirically, we map the theoretical objects $\{s_{\text{imp}}^{\text{GDP}}, w_e, \Delta \ln \kappa_e\}$ to observable data. The import share in GDP, $s_{\text{imp}}^{\text{GDP}}$, is obtained directly from national accounts. The lane weight w_e is constructed from observed TEU volumes,

$$w_e = \text{TEU}_e \big/ \sum_{e'} \text{TEU}_{e'},$$

ensuring that heavily used routes receive greater weight in the welfare calculation.

The remaining object is the route-level cost change $\Delta \ln \kappa_e$ implied by the surcharges. In the model, a surcharge on fleet n raises its rental rate by the ad valorem amount T_n, which increases the iceberg cost on any lane that uses that fleet in proportion to its exposure:

$$\Delta \ln \kappa_e \approx T_n \, a_{e,n} \, \frac{p_e}{p_e + v_e}.$$

To operationalize this expression, we convert the statutory per-ton and per-container fees into their ad valorem equivalents $T_{b,t}$ for each vessel size class b in period t. Let $s_{\text{CHN}}(e, b)$ denote the TEU-weighted share of capacity on lane e provided by the targeted fleets in size class b, and let $s_{\text{ship}}^{\text{imp}}$ denote the share of shipping costs in importer expenditure. The effective ad valorem surcharge applied to shipments on

lane e in period t is then

$$\tau_{e,t}^{\text{eff}} = \sum_{b} T_{b,t}\, s_{\text{CHN}}(e,b)\, s_{\text{ship}}^{\text{imp}}.$$

This expression combines the policy shock, the composition of vessels serving each route, and the contribution of shipping charges to delivered import prices.

Putting these elements together, the welfare effect in period t is approximated by

$$\frac{\Delta \text{EV}_t}{Y} \approx - s_{\text{imp}}^{\text{GDP}} \sum_{e} w_e\, \tau_{e,t}^{\text{eff}}.$$

This construction highlights the role of lane-level exposure: routes that make intensive use of the targeted fleets experience larger increases in effective shipping costs, while routes that rarely use those fleets are largely unaffected. Because imports represent a modest share of U.S. GDP, the aggregate effect of the surcharges remains small even when the route-level cost increases are sizable.

6.4 Quantitative Results

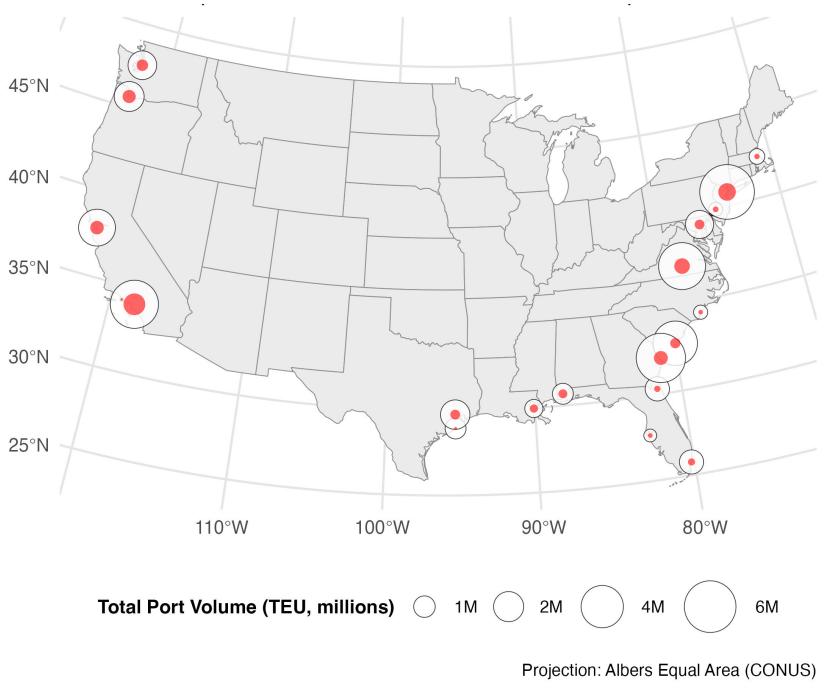
The welfare expression derived in the previous subsection indicates that the aggregate impact of the surcharges depends on two empirical components: the effective surcharge $\tau_{e,t}^{\text{eff}}$ applied to each lane and the economic importance of each lane, captured by the weights w_e. Both objects vary substantially across routes, because U.S. shipping lanes differ sharply in how intensively they rely on the fleets subject to the policy. We therefore begin by documenting the underlying exposure patterns that shape the incidence of the surcharges.

Because the surcharges apply to specific classes of ships, the effect on any particular route depends on how intensively that route relies on the taxed fleets. We define the unit of analysis as a directed shipping lane $e = (o \to d)$ and a ship-size category $b \in \{< 3{,}000,\ 3{,}000{-}8{,}000,\ 8{,}000{-}12{,}000,\ 12{,}000{-}17{,}000\}$ TEUs. For each lane–size pair, we compute $s_{\text{CHN}}(e,b)$, the TEU-weighted share of capacity carried on Chinese-built ships. These exposure shares are key determinants of the effective surcharge on each lane, since lanes that rely disproportionately on the targeted fleets experience larger increases in shipping costs.

To illustrate how these exposure patterns vary across the country, Figure 9 maps the TEU-weighted share of U.S. imports carried on Chinese-built ships by port. The geographic distribution mirrors the lane-level patterns documented above: West Coast ports, particularly those handling trans-Pacific traffic, exhibit the highest dependence on Chinese-built capacity, while ports on the East Coast and Gulf Coast rely far less on these fleets. This spatial heterogeneity reinforces the importance of measuring exposure at the lane level, as the incidence of fleet-specific surcharges varies sharply across regions and shipping corridors.

We now combine these exposure patterns with the surcharge schedule to quantify the welfare effects. Our analysis evaluates the impact of the fee schedule as originally proposed for implementation beginning in October 2025, prior to the subsequent postponement of the policy to late 2026. Table 1 reports the results of aggregating lane-level shocks using the welfare mapping in Section 4. As the statutory fees rise from 0.20 to 0.35 in ad valorem terms between 2025 and 2028, the national equivalent-variation

Figure 9. U.S. Import Dependence on Chinese-Built Containerships by Port



Notes: The figure reports the TEU-weighted share of U.S. imports carried on Chinese-built containerships by port, using the merged Panjiva–Clarksons dataset. West Coast ports exhibit the highest exposure, reflecting their central role in trans-Pacific shipping networks.

Table 1. Estimated Welfare Effects of U.S. Fees on Chinese Ships

Date	Ad valorem fee (%)	EV ($ bn)	EV / GDP (%)	Effective tariff (%)
October 2025	0.20	3.43	0.01	0.05
April 2026	0.25	4.44	0.02	0.06
April 2027	0.30	5.31	0.02	0.07
April 2028	0.35	6.17	0.02	0.08

Notes: Equivalent variation (EV) represents welfare losses from U.S. Section 301 surcharges on Chinese-built and Chinese-operated containerships, aggregated across all shipping lanes. Ad valorem rates correspond to effective per-container equivalents of the statutory fees, computed by ship size. Welfare-to-GDP ratios use 2025 U.S. nominal GDP.

(EV) loss increases from $3.4 billion to $6.2 billion. The implied welfare losses range from 0.01 to 0.02 percent of U.S. GDP.

Although the aggregate effects are modest, they are far from uniform across shipping routes. Lanes connecting Asia to the U.S. West Coast—especially those served disproportionately by medium-sized (3,000–8,000 TEU) Chinese-built ships—bear the bulk of the burden. These routes face both higher statutory surcharges and higher exposure shares $s_{\mathrm{CHN}}(e, b)$, which combine to generate elevated effective surcharges $\tau^{\mathrm{eff}}_{e,t}$. In contrast, East Coast and Europe–U.S. lanes are only lightly affected, reflecting their lower reliance on the targeted fleets and the availability of alternative capacity. These patterns mirror the heterogeneous propagation emphasized in Section 4: direct exposure, network position, and substitution possibilities jointly shape the incidence of shipping-specific policy shocks.

The quantitative evaluation indicates that the surcharges impose small but non-negligible welfare losses on the United States. The limited aggregate magnitude reflects two structural features. First, shipping costs constitute only a modest share of import values, so even large percentage increases in shipping costs translate into muted changes in delivered import prices. Second, U.S. import routes are served by a globally diversified fleet with substantial scope for substitution across vessels and corridors, which dampens the aggregate effect of policies targeting specific fleets. At the same time, the concentration of losses on a subset of trans-Pacific lanes highlights how targeted interventions in logistics can generate highly uneven regional impacts, even when aggregate effects are small.

More broadly, the exercise demonstrates how policies aimed at specific shipping fleets propagate through the global network of trade routes and affect the cost and structure of U.S. imports. Even in the absence of changes to tariffs on goods, interventions in the maritime logistics sector can reshape delivered trade costs, alter routing patterns, and influence welfare. As concerns about the reliability of global supply chains continue to rise, understanding these channels will become increasingly important for the design and evaluation of future policies.

7 Quantifying Vulnerabilities to Shipping Disruptions

This section quantifies the vulnerability of the U.S. economy to a disruption in access to Chinese-built containerships. While the previous section evaluates policies designed to alter fleet composition, here we consider the counterfactual scenario in which access to Chinese-built capacity serving U.S. trade is restricted. This exercise provides a direct measure of how concentration in global shipbuilding translates into trade cost increases and welfare losses in the event of a geopolitical disruption.

We proceed in five steps. First, we describe the parametrization of the model. Second, we define the counterfactual disruption that removes access to Chinese-built fleets. Third, we explain how the model is solved under this counterfactual. Fourth, we report the quantitative results. Finally, we discuss the mechanisms that govern the magnitude of the welfare effects.

7.1 Parametrization

The model is calibrated to match key features of global trade, port-level activity, and shipping markets. The unit of analysis in the quantitative implementation is the port. Each port is treated as a node in the shipping network, and trade flows occur between origin and destination ports. Table 2 summarizes the parameter values and their sources.

Predetermined Parameters. The elasticity of substitution across goods, σ, governs trade responsiveness to changes in delivered prices and follows standard estimates in the quantitative trade literature. The congestion parameter δ controls how shipping costs scale with traffic on a given edge. Setting $\delta = 1$ implies constant returns to volume in shipping services, providing a benchmark in which congestion effects operate proportionally to traffic.

Parameter	Description	Value	Target / Source
Predetermined			
σ	Elasticity of substitution over goods	4	Simonovska and Waugh (2014)
δ	Congestion parameter	1	Constant returns to volume
η	Fleet substitution elasticity	3.0	High substitutability baseline
ψ	Weight on fleet vs. non-fleet transport	0.5	Symmetric dependence
ω_{Total}	Overall trade openness	0.4	Match tradable share of global GDP
Estimated			
A_i, u_i	Location productivity and amenity		Match GDP and population (wL_i, L_i)
τ_{ij}	Bilateral trade frictions		Match bilateral trade shares π_{ij}
κ_{ij}^f	Fleet-specific transport costs		Match fleet market shares $s_{n,e}$

Table 2. Parameter Values and Sources

The elasticity of substitution across fleets, η, determines the scope for substitution between Chinese-built vessels and alternative shipbuilding origins. A higher η implies that operators can more easily reallocate capacity across fleets when access to a particular origin is disrupted. The baseline value of $\eta = 3$ reflects a setting with substantial, but not perfect, substitutability.

The parameter ψ governs the relative importance of shipping costs in total delivered trade costs. It determines how changes in shipping prices translate into iceberg trade costs. Finally, the trade openness parameter ω_{Total} is set to match the observed share of tradable output in global GDP, anchoring the model's aggregate exposure to international trade.

Port-Level Economic Fundamentals. The quantitative implementation is conducted at the port level. Since GDP and population data are available at the country level from the World Bank, we allocate national income and population to ports using weights derived from shipping activity. Specifically, for each country, we compute the fraction of arriving containership traffic handled by each port based on observed port calls and TEU volumes. National GDP and population are then distributed across ports proportionally to these weights.

Formally, if country m has GDP Y_m and population L_m, and port $i \in m$ accounts for a share ω_i^m of the country's shipping traffic, then port-level output and labor are assigned as

$$Y_i = \omega_i^m Y_m, \quad L_i = \omega_i^m L_m.$$

This procedure ensures that port-level economic size reflects observed integration into global shipping networks, while preserving consistency with national accounts.

Location-specific productivities A_i and amenities u_i are then calibrated so that the model reproduces these assigned port-level GDP and population levels in equilibrium. This approach allows the model to match both the geographic distribution of economic activity and observed shipping flows.

Trade Frictions and Fleet Costs. Bilateral trade frictions τ_{ij} are disciplined to match observed bilateral trade shares across ports. Using exact hat algebra, we recover the trade cost structure consistent with

observed trade patterns and the chosen elasticity σ.

Fleet-specific transport cost components κ_{ij}^{f} are calibrated to match observed fleet market shares across lanes. For each directed port-to-port lane, we observe the share of capacity supplied by different fleet origins. The model's CES aggregation of fleets implies cost shares that depend on relative transport costs. We therefore choose κ_{ij}^{f} so that the model reproduces these observed fleet market shares in the baseline equilibrium.

Together, these calibrations ensure that the baseline model replicates three key features of the data: the geographic distribution of economic activity, observed bilateral trade flows, and the allocation of shipping capacity across fleet origins.

This baseline provides the foundation for the counterfactual experiments that follow, in which access to specific fleets is restricted and the resulting adjustments in shipping prices, routing patterns, trade flows, and welfare are computed in general equilibrium.

7.2 Counterfactual: Disruption to Chinese-Built Fleets

Rising geopolitical tensions between the United States and China have increased scrutiny of dependencies in critical transport infrastructure. In maritime shipping, these concerns are salient because global shipbuilding is highly concentrated in Chinese and South Korean shipyards, and Chinese-built vessels account for a significant share of the capacity serving U.S. trade. A disruption to this capacity would therefore have direct implications for shipping costs and trade flows.

To quantify this exposure, we consider a deterministic counterfactual in which the United States loses access to Chinese-built containerships on routes serving U.S. ports. The objective is to measure how dependent current trade patterns are on this concentrated source of shipping capacity. Comparing the baseline equilibrium to one in which Chinese-built fleets are unavailable provides a benchmark measure of vulnerability.

Formally, for fleets indexed by $n \in \mathcal{F}_{CHN}$, we set the availability indicator to zero on all edges serving U.S. ports:

$$\Omega_{e,n} = 0 \quad \text{for all } e \in \mathcal{E}_{US},\ n \in \mathcal{F}_{CHN}.$$

All other fleets remain available, and trade policy on goods is unchanged. The disruption operates exclusively through the shipping sector.

Because the model is solved in a deterministic setting, the resulting welfare changes capture the difference between two steady states: one with full access to Chinese-built fleets and one without such access. This difference measures the structural vulnerability implied by concentration in global shipbuilding.

7.3 Implementation

We compute the counterfactual equilibrium using exact hat algebra. This approach allows us to solve for proportional changes in wages, shipping costs, trade costs, and welfare relative to the calibrated baseline, without re-estimating levels of productivities or bilateral trade frictions. The baseline equilibrium already matches observed port-level income and population (by construction), bilateral trade patterns, and

the allocation of shipping capacity across fleets. The counterfactual then modifies only the relevant primitives and recomputes the equilibrium adjustments.

The disruption to Chinese-built fleets is implemented as a change in fleet availability on U.S.-serving routes. For all U.S.-bound lanes, we remove Chinese-built vessels from the set of fleets that can supply shipping services. This affects shipping costs directly on exposed lanes, and indirectly through changes in fleet rental rates and traffic reallocation across routes.

The solution proceeds as a nested fixed point with two tightly linked blocks. The first block solves the real side of the economy given trade costs: wages and nominal incomes adjust across ports so that goods markets clear and labor reallocates within countries. The second block solves the shipping sector given traffic: operators substitute across available fleets, fleet rental rates adjust to clear capacity markets, and lane-level shipping prices update given congestion and the composition of fleets used on each lane. These updated lane-level shipping prices imply new generalized trade costs, which feed back into the real side by changing delivered prices and trade flows. We iterate between these two blocks until lane-level trade costs, wages, and traffic are jointly consistent.

This procedure yields an equilibrium in which the loss of Chinese-built capacity is absorbed through three channels. First, lanes that were heavily reliant on Chinese-built vessels experience higher shipping costs. Second, operators substitute toward alternative fleets and, where feasible, toward alternative routes, reassigning traffic across the network. Third, these shipping adjustments feed back into trade flows and relative prices, generating general equilibrium changes in wages, expenditures, and welfare. The counterfactual welfare effect is then computed by comparing the implied welfare index for the United States in the disruption equilibrium to the baseline.

7.4 Quantitative Results (preliminary)

We begin by examining the effects of the disruption on lane-level shipping costs. The loss of access to Chinese-built fleets generates substantial heterogeneity across routes. Lanes that rely heavily on Chinese-built vessels experience the largest increases in generalized shipping costs, while routes with limited baseline exposure are only modestly affected. In particular, trans-Pacific routes serving the U.S. West Coast exhibit the strongest responses, reflecting their high reliance on medium-sized Chinese-built vessels in the baseline equilibrium. By contrast, many East Coast and Europe–U.S. routes adjust primarily through substitution toward alternative fleets and therefore experience smaller cost increases.

The disruption also affects fleet rental markets. Removing Chinese-built capacity increases demand for alternative fleets, raising their rental rates and spreading the shock beyond directly exposed lanes. This rental-rate channel amplifies the initial impact by transmitting the disruption to routes that do not directly rely on Chinese-built vessels. At the same time, substitution across fleets and corridors mitigates the overall increase in shipping costs. The net effect depends on the elasticity of substitution across fleets and on the distribution of exposure across lanes.

At the aggregate level, the disruption leads to an increase in delivered trade costs and a decline in welfare. Although shipping costs represent a relatively modest share of delivered prices, the concentration of exposure on major trade corridors implies that the disruption produces measurable welfare losses. The magnitude of the decline depends on three structural features emphasized by the model: the baseline

exposure to Chinese-built capacity, the elasticity of substitution across fleets, and the share of shipping costs in total import expenditure.

Beyond the aggregate effect, the disruption produces meaningful reallocation across ports and routes. Traffic shifts toward lanes served by alternative fleets, and relative wages adjust across ports in response to changes in trade costs. These adjustments highlight that vulnerability is not uniform across the economy: regions that are more dependent on trans-Pacific shipping bear a disproportionate share of the burden.

Taken together, the results quantify the structural vulnerability created by concentration in global shipbuilding. Access to Chinese-built fleets lowers shipping costs in the baseline equilibrium, but it also exposes trade flows to disruptions that propagate through shipping markets and trade networks. The counterfactual therefore provides a benchmark measure of exposure that informs the evaluation of policies aimed at reducing reliance on concentrated foreign fleets.

8 Insights from a Stylized Numerical Exercise

This section presents a stylized numerical exercise that complements the port-level quantitative analysis in the previous section. The goal of the exercise is to isolate the core mechanisms of the model in a simplified environment, allowing us to examine more directly how structural trade dependence and substitution across fleets shape the transmission of a fleet-specific disruption into trade costs and welfare.

The central question of the paper is how dependence on a concentrated foreign shipping supply base shapes the efficiency and vulnerability of trade, and how fleet-specific disruptions or interventions propagate through shipping markets and trade networks. The full quantitative results in the main text answer this question using global fleet data and detailed exposure patterns. The stylized exercise here isolates the same forces in a minimal environment. In particular, it highlights two general determinants of vulnerability: (i) baseline dependence on international trade and shipping, and (ii) the ability of shipping services to substitute across fleets when one fleet becomes more costly.

8.1 Environment

The stylized economy is a directed network with five nodes partitioned into two countries. Country 1 is a "major economy" consisting of three nodes, indexed $i \in \{1, 2, 3\}$, and is relatively internally oriented. Country 2 is a "partner economy" consisting of two nodes, indexed $i \in \{4, 5\}$, and is more reliant on international trade. Nodes are connected by directed edges representing shipping lanes.

Shipping services on each edge are assembled from fleet capacity using a CES aggregator with elasticity of substitution $\eta > 1$. We consider three fleets: (i) a domestic fleet that serves only within-country edges, (ii) an international "safe" fleet that serves cross-country edges and is unaffected by the shock, and (iii) an international "target" fleet that serves cross-country edges and is subject to a fleet-specific disruption. Fleet rental rates adjust endogenously to clear fleet markets. Shipping prices on each edge reflect the composition of fleets used on that route, and generalized iceberg trade costs κ_e are determined in equilibrium.

8.2 Parametrization and Baseline Dependence

The numerical exercise is not intended as a calibrated representation of the global shipping system. Instead, it fixes a small set of standard elasticities and chooses baseline lane weights and fleet shares to generate an asymmetric pattern of import dependence across the two countries. We set the goods substitution elasticity to $\sigma = 4$, consistent with the values used in the main analysis. We set the fleet substitution elasticity to $\eta = 1.5$, corresponding to a relatively limited but nontrivial ability to substitute across fleets. This parameterization is useful for illustrating how imperfect substitution amplifies differential exposure across countries.

The baseline equilibrium implies markedly different degrees of import dependence across the two countries. We measure a country's import dependence as the fraction of its absorption that is sourced using cross-border edges in the baseline equilibrium. Under the baseline topology and lane weights, the major economy has import dependence 0.12, while the partner economy has import dependence 0.48. These asymmetries are embedded in the network structure and provide the key ingredient for heterogeneous welfare effects in the counterfactual.

The network topology underlying this asymmetry is illustrated in Panel A of Figure 10. The major economy has substantial internal connectivity, and a large share of its absorption can be served through domestic routes. The partner economy, in contrast, is structurally more exposed to cross-border lanes and therefore to the international fleets that supply them.

8.3 Counterfactual Shock

We study a fleet-specific disruption affecting one of the international fleets. In particular, we introduce a cost increase for the international "target" fleet while leaving all other fleets unchanged. Formally, let μ_n denote the cost shifter governing the supply of capacity from fleet n. Relative to the baseline, we set

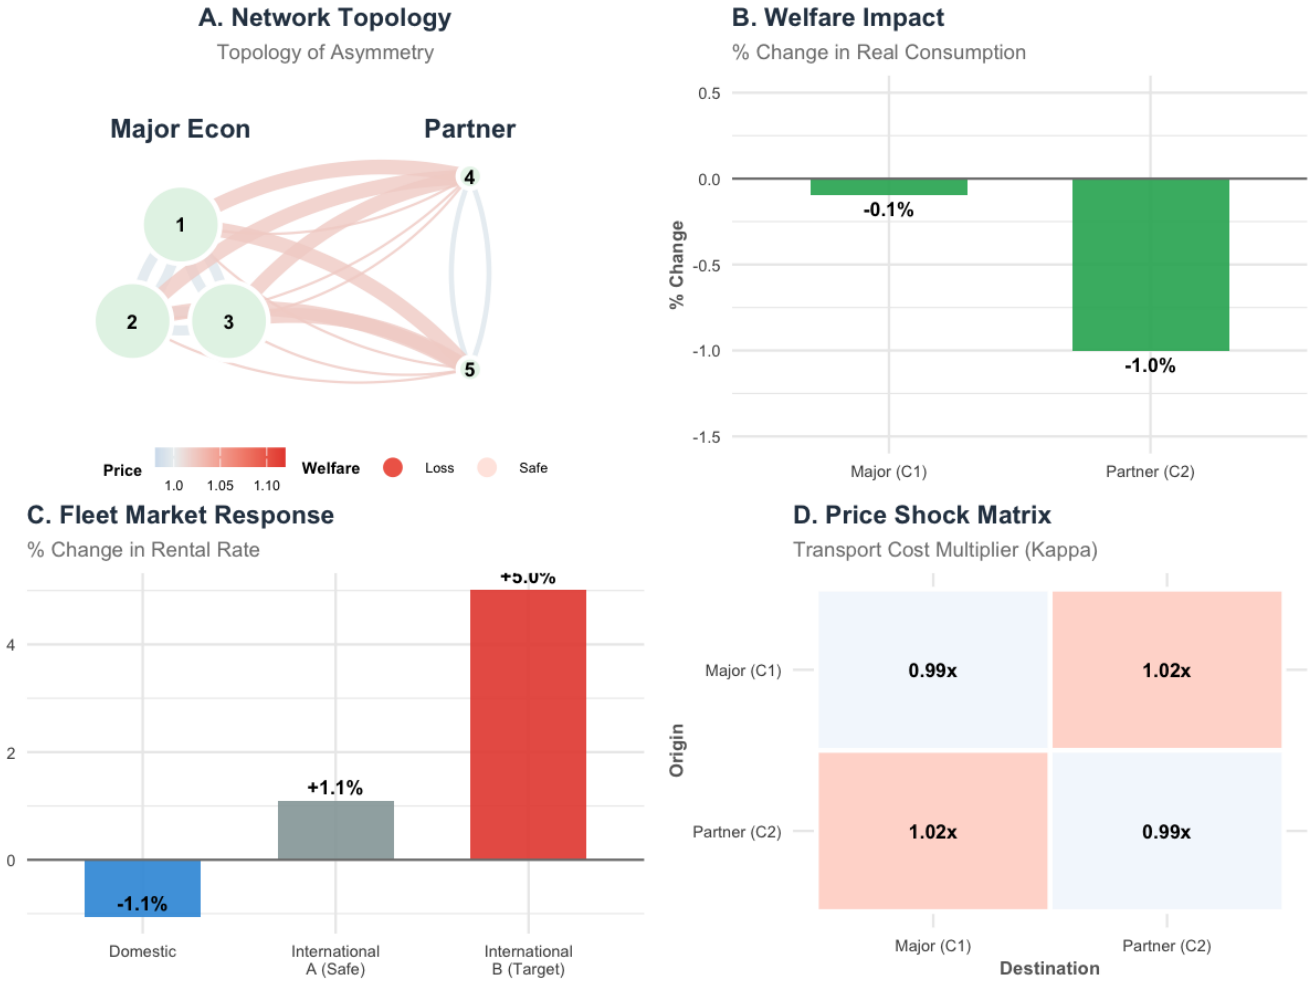
$$\hat{\mu}_{\text{Domestic}} = 1, \qquad \hat{\mu}_{\text{Intl Safe}} = 1, \qquad \hat{\mu}_{\text{Intl Target}} = 1.1.$$

This specification corresponds to a 10 percent increase in the effective cost of supplying capacity from the targeted fleet.

All other primitives are held fixed. The counterfactual therefore isolates how an increase in the cost of a single fleet propagates through fleet rental markets, alters the composition of shipping services across routes, and ultimately affects trade costs and welfare in general equilibrium.

We solve the counterfactual equilibrium using the same hat-algebra implementation strategy as in the main quantitative analysis. Starting from the baseline equilibrium, the counterfactual modifies only the fleet cost shifter for the targeted fleet and recomputes a fixed point in wages, nominal incomes, lane-level generalized costs, and fleet rental rates. Equilibrium adjustments occur through three channels that mirror the main model: direct exposure of affected lanes, substitution across fleets, and general equilibrium adjustments in trade flows and prices.

Figure 10. Stylized Numerical Exercise



Notes: Panel A plots the five-node topology with asymmetric reliance on international connections. Panel B reports equilibrium welfare changes by country. Panel C reports equilibrium changes in fleet rental rates. Panel D reports the country-to-country matrix of trade cost multipliers implied by the disruption.

8.4 Results

Lane-level trade cost adjustments. The disruption generates sharp heterogeneity across lanes. Because the targeted fleet is used disproportionately on cross-border edges, generalized trade costs rise on international lanes and fall slightly on domestic lanes, reflecting reallocation and rental-rate adjustments. Panel D of Figure 10 summarizes these changes in a compact "price shock matrix" at the country-to-country level. The implied trade cost multipliers are approximately 1.02 on cross-border routes and 0.99 on within-country routes.

Table 3 reports the corresponding weighted average cost multipliers, computed from lane-level $\hat{\kappa}_e$ using baseline lane weights. The key feature is that the disruption increases the cost of cross-border trade in both directions, while within-country trade becomes slightly cheaper as the equilibrium shifts away from international shipping and into domestic sourcing where feasible.

Table 3. Implied Trade Cost Multipliers by Origin and Destination Country ($\hat{\kappa}$)

	Destination: Major (C1)	**Destination: Partner (C2)**
Origin: Major (C1)	0.993	1.022
Origin: Partner (C2)	1.022	0.993

Table 4. Fleet Market Impacts

Fleet	Rental Rate (% Δ)	Market Share (Baseline)	Cost Shock
Domestic	-1.07%	0.54	1.0x
International A (Safe)	1.09%	0.23	1.0x
International B (Target)	5.02%	0.23	1.1x

Fleet market response and substitution. A central feature of the model is that fleet rental rates respond endogenously, dampening or amplifying the incidence of fleet-specific shocks. Panel C of Figure 10 reports the equilibrium rental-rate changes across fleets. The rental rate of the targeted international fleet rises by about 5 percent, while the rental rate of the safe international fleet rises by about 1 percent. The domestic fleet rental rate falls by about 1 percent, reflecting reduced demand for international capacity and reallocation toward domestic sourcing on internal edges.

Table 4 reports these changes together with baseline fleet market shares. The baseline shares are roughly split between the domestic fleet and the two international fleets, with each international fleet accounting for about 23 percent of baseline activity. The fact that a 10 percent cost increase for the targeted fleet produces a smaller than proportional increase in its equilibrium rental rate illustrates the importance of substitution. Operators shift demand toward the safe fleet, and the general equilibrium shift in lane traffic reallocates shipping services away from international routes where possible.

Welfare effects and the role of trade dependence. Panel B of Figure 10 reports welfare effects, measured as the percent change in real consumption. The major economy experiences a small welfare loss of about 0.09 percent, while the partner economy experiences a much larger loss of about 1.0 percent. The difference reflects the structural asymmetry built into the baseline network: the partner economy is substantially more trade-dependent and therefore more exposed to increases in international shipping costs.

Table 5 decomposes these welfare changes into movements in the price index and nominal wages. In this exercise, nominal wages are essentially unchanged, and welfare effects operate almost entirely through changes in the consumption price index. This price-index channel is the direct analogue of the mechanism emphasized in the main quantitative analysis. When a fleet-specific disruption raises shipping costs on the routes that connect a country to its suppliers, delivered prices rise, and welfare falls in proportion to the importance of trade in absorption and the degree of route exposure.

Discussion. Although the environment here is intentionally stylized, it highlights three general lessons that carry over to the port-level quantitative exercises.

Table 5. Simulation Results: Asymmetric Fleet Shock (% Change)

Country	Real Welfare	Price Index	Nominal Wage	Import Dependence	Hulten Prediction
Major Economy (C1)	-0.09%	0.09%	0.00%	0.12	0.00%
Partner Economy (C2)	-1.00%	1.01%	0.00%	0.48	0.00%

First, vulnerability is fundamentally shaped by baseline dependence on trade and on the shipping network. The partner economy experiences a much larger welfare loss because a larger fraction of its absorption is tied to cross-border edges that become more expensive when the targeted fleet is disrupted.

Second, fleet substitution dampens the mapping from a fleet-specific shock to equilibrium rental rates and shipping prices. The shock raises the rental rate of the targeted fleet, but by less than the underlying cost change, as operators substitute toward the safe fleet and as traffic reallocates across lanes. This mechanism provides intuition for why the incidence of fleet-targeted policies in the main quantitative analysis depends not only on exposure, but also on substitution possibilities and market clearing in fleet rental markets.

Third, welfare changes operate primarily through delivered prices rather than through large movements in nominal income. In the model, shipping is an input into market access. When a disruption raises generalized trade costs on the relevant edges, the price index rises and real consumption falls. This price-index channel is central to interpreting both the policy experiment and the vulnerability counterfactual in the empirical sections of the paper.

The stylized exercise highlights the central tension studied in the paper. Concentration in shipping capacity delivers efficiency in normal times, but it creates exposure when access to a particular fleet is disrupted. The magnitude of the resulting welfare loss depends on structural trade dependence and on the scope for substitution across fleets. These same forces underlie the empirical results on U.S. exposure to Chinese-built vessels.

9 Concluding Remarks

This paper documents the structure of the global containership fleet and quantifies how concentration in shipbuilding and fleet ownership shapes trade costs and exposure in international shipping. Using vessel-level data from the global registry and shipment-level trade data, we show that global shipbuilding is highly concentrated in a small number of East Asian economies, while vessel operation is considerably more dispersed. Countries such as the United States rely heavily on this concentrated supply base, particularly on Chinese-built capacity in specific segments of the fleet.

We embed these empirical patterns in a tractable general equilibrium framework that links fleet composition to shipping prices, routing decisions, trade costs, and welfare. The model allows us to quantify both the vulnerability implied by concentrated shipbuilding and the effects of policies that seek to alter fleet composition. Counterfactual disruptions to Chinese-built capacity generate heterogeneous increases in shipping costs across routes, reflecting exposure and substitution possibilities within the

global network. Policies that reduce reliance on particular fleets mitigate this exposure but do so at the cost of higher trade costs and lower welfare.

Taken together, the evidence and the framework highlight how the organization of global shipping influences both the efficiency and the resilience of international trade. Concentration in upstream transportation infrastructure delivers scale economies and lower costs in normal times, yet it also creates exposure to geopolitical and supply-chain disruptions. By providing a disciplined mapping from fleet-level shocks to trade and welfare outcomes, this paper contributes to a broader understanding of how strategic dependence in critical infrastructure shapes the transmission of economic shocks across borders.

References

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.

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.

Baldwin, Richard. 2022. "Globotics and Macroeconomics: Globalisation and Automation of the Service Sector." National Bureau of Economic Research NBER Working Paper 30317.

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*, rdaf011.

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

Bonadio, Barthélémy. 2021. "Ports vs. Roads: Infrastructure, Market Access and Regional Outcomes." Working Paper.

Bonadio, Barthélémy, Zhen Huo, Andrei A. Levchenko, and Nitya Pandalai-Nayar. 2021. "Global Supply Chains in the Pandemic." *Journal of International Economics*, 133: 103518.

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

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

Coşar, A. Kerem, and Banu Demir. 2016. "Domestic Road Infrastructure and International Trade: Evidence from Turkey." *Journal of Development Economics*, 118: 232–244.

Coşar, A. Kerem, and Pablo D. Fajgelbaum. 2016. "Internal Geography, International Trade, and Regional Specialization." *American Economic Journal: Microeconomics*, 8(1): 24–56.

Cristea, Anca, David Hummels, Laura Puzzello, and Misak Avetisyan. 2013. "Trade and the greenhouse gas emissions from international freight transport." *Journal of environmental economics and management*, 65(1): 153–173.

Dornbusch, R, S Fischer, and P A Samuelson. 1977. "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods." *Am. Econ. Rev.*, 67(5): 823–839.

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

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.

Fan, Jingting, Yi Lu, and Wenlan Luo. 2023. "Valuing domestic transport infrastructure: a view from the route choice of exporters." *The Review of Economics and Statistics*, 1–46.

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.

Freund, Caroline, Maryla Maliszewska, Aaditya Mattoo, and Michele Ruta. 2020. "When Elephants Make Peace: The Impact of the China–U.S. Trade Agreement on Developing Countries." World Bank Policy Research Working Paper 9173.

Fuchs, Simon, and Woan Foong Wong. 2024. "Multimodal Transport Networks." CEPR Discussion Paper 19531. CEPR Press, Paris & London.

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.

Gawande, Kishore, Pravin Krishna, and Marcelo Olarreaga. 2009. "What Governments Maximize and Why: The View from Trade." *International Organization*, 63(3): 491–532.

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.

Heiland, Inga, Andreas Moxnes, Karen-Helene Ulltveit-Moe, and Yuan Zi. 2019. "Trade from Space: Shipping Networks and the Global Implications of Local Shocks." Centre for Economic Policy Research CEPR Discussion Paper 14193.

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

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

Kreindler, Gabriel E, and Yuhei Miyauchi. 2023. "Measuring commuting and economic activity inside cities with cell phone records." *Review of Economics and Statistics*, 105(4): 899–909.

Lugovskyy, Volodymyr, Alexandre Skiba, and David Terner. 2025. "Unintended consequences of environmental regulation of maritime shipping: Carbon leakage to air shipping." *Journal of International Economics*, 155: 104081.

Miyauchi, Yuhei, Kentaro Nakajima, and Stephen J Redding. 2025. "The economics of spatial mobility: Theory and evidence using smartphone data." *The Quarterly Journal of Economics*, 140(4): 2507–2570.

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.

Rodrigue, Jean-Paul, Theo Notteboom, and Athanasios A. Pallis. 2011. "The Financialization of the Port and Terminal Industry: Revisiting Risk and Embeddedness." *Maritime Policy & Management*, 38(2): 191–213.

Rodrik, Dani. 2022. "An Industrial Policy for Good Jobs." The Hamilton Project Policy Proposal.

Rodrik, Dani, and Charles F. Sabel. 2019. "Building a Good Jobs Economy." In *Good Jobs for All in a Changing World of Work*. Oxford University Press.

Shapiro, Joseph S. 2016. "Trade Costs, CO2, and the Environment." *American Economic Journal: Economic Policy*, 8(4): 220–54.

Tsivanidis, Nick. 2022. "Evaluating the impact of urban transit infrastructure: Evidence from bogota's transmilenio." *Unpublished manuscript*, 18.

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

Appendix

This appendix is organized in three parts. The first develops a stylized benchmark that links route allocation, endogenous transport wedges, DFS cutoffs, stochastic capacity, and optimal policy. The second records the quantitative network environment, including shipping technology, soft-min routing, pass-through, welfare aggregation, risk adjustment, and equilibrium. The third presents the traffic-based hat-algebra system and the numerical procedure used for fleet-specific counterfactuals.

A Simple Model: DFS with endogenous routing

This section develops a tractable benchmark that isolates the core mechanisms emphasized in the paper. We begin from a standard DFS production environment, introduce an endogenous route-allocation problem that generates transport wedges, embed those wedges into the goods cutoffs, and then extend the setup to stochastic capacity and policy analysis.

A.1 Setup, notation, and roadmap

Roadmap. Subsection A.2 presents the standard DFS production side. Subsection A.3 derives the deterministic route-allocation problem and clarifies how the active support is determined. Subsection A.4 gives the closed-form constant-capacity benchmark together with the interior-support condition and the exact capacity elasticity of transport costs. Subsection A.5 embeds the endogenous transport wedge into the DFS cutoffs. Subsection A.6 introduces stochastic capacity, states the required moment conditions, and reports the reliability-premium approximation. Subsection A.7 defines the macro elasticity objects used in policy analysis. Subsection A.8 derives the broad policy rule, the exact scalar rule under a uniform capacity subsidy, and the factorized approximation used in the summary.

A.2 The DFS production benchmark

Technology and ordering. There are two countries and a continuum of goods $g \in [0,1]$. Unit labor requirements are $a_H(g)$ and $a_F(g)$, wages are $w_H, w_F > 0$, and

$$A(g) := \ln \frac{a_F(g)}{a_H(g)}.$$

Assumption A.1 (DFS ordering)**.** The relative productivity schedule is strictly decreasing:

$$A'(g) < 0 \qquad \forall g \in [0,1].$$

Let $\varpi = \ln(w_H/w_F)$. In the frictionless benchmark, Home supplies good g whenever $A(g) \geq \varpi$, and the unique frictionless cutoff g_0 solves

$$A(g_0) = \varpi.$$

A.3 Endogenous routing without risk

Route space and transport problem. For an imported good g, route choice is over $x \in [0,1]$. Let $f_g(x) \geq 0$ denote route density, with total volume

$$\int_0^1 f_g(x)\,\mathrm{d}x = m_g.$$

Transport cost is

$$m_g T_g = \min_{f_g(\cdot)\geq 0} \left\{ \int_0^1 \left[\tau_0 + \beta_g\Big(x - \tfrac{1}{2}\Big)^2\right] f_g(x)\,\mathrm{d}x + \frac{\gamma}{2}\int_0^1 \frac{f_g(x)^2}{K(x)}\,\mathrm{d}x \right\}.$$

First-order condition. Let μ_g be the multiplier on the flow constraint. Then

$$\tau_0 + \beta_g\Big(x - \tfrac{1}{2}\Big)^2 + \gamma\frac{f_g(x)}{K(x)} \geq \mu_g$$

with equality if $f_g(x) > 0$, so

$$f_g^\star(x) = \frac{K(x)}{\gamma}\left[\mu_g - \tau_0 - \beta_g\Big(x - \tfrac{1}{2}\Big)^2\right]_+. \tag{A.1}$$

Proposition A.1 (Route-support cutoff)**.** *Suppose $K(x) > 0$ for all $x \in [0,1]$. Then the active route set of good g is*

$$\mathcal{X}_g^\star = \left[\tfrac{1}{2} - \Delta_g, \tfrac{1}{2} + \Delta_g\right]$$

for some $\Delta_g \in (0, \frac{1}{2}]$.

Proof. From (A.1), route x is used if and only if

$$\mu_g - \tau_0 - \beta_g\Big(x - \tfrac{1}{2}\Big)^2 > 0.$$

Since $K(x)$ is strictly positive, it scales the optimal density but does not affect the positivity condition. The quadratic term is symmetric and strictly increasing in $|x - \frac{1}{2}|$, so the active set is an interval centered at the midpoint. □

Remark A.1 (What capacity heterogeneity can and cannot do). Under the present specification, local capacity heterogeneity changes the density of traffic within the active set but does not by itself create holes in route support. To obtain disconnected route usage sets one would need literal outages, zero-capacity segments, or additional route-specific fixed costs.

Distance-from-midpoint representation. Let $u = |x - \frac{1}{2}| \in [0, \frac{1}{2}]$, and write $K(u)$ for symmetric capacity. At the edge of the active set, $f_g^\star(\Delta_g) = 0$, so $\mu_g = \tau_0 + \beta_g\Delta_g^2$. Substituting gives

$$f_g^\star(u) = \frac{\beta_g K(u)}{\gamma}\Big(\Delta_g^2 - u^2\Big), \qquad 0 \leq u \leq \Delta_g.$$

Integrating yields the shipment constraint

$$m_g = \frac{2\beta_g}{\gamma}\int_0^{\Delta_g} K(u)\Big(\Delta_g^2 - u^2\Big)\,\mathrm{d}u, \tag{A.2}$$

and substituting back into the objective gives the per-unit transport cost

$$T_g = \tau_0 + \frac{\beta_g^2}{m_g \gamma} \int_0^{\Delta_g} K(u)\Big(\Delta_g^4 - u^4\Big)\,\mathrm{d}u. \tag{A.3}$$

A.4 Closed-form benchmark with constant capacity

Suppose

$$K(x) \equiv K_0.$$

Assumption A.2 (Interior support under constant capacity)**.**

$$m_g < \frac{K_0 \beta_g}{6\gamma}.$$

Proposition A.2 (Closed-form route width and transport cost)**.** *Under constant capacity and Assumption A.2,*

$$\Delta_g = \left(\frac{3\gamma m_g}{4K_0\beta_g}\right)^{1/3},$$

and

$$T_g = \tau_0 + C_0 \left(\frac{\gamma}{K_0}\right)^{2/3} \beta_g^{1/3} m_g^{2/3}, \qquad C_0 := \frac{3}{5}\left(\frac{3}{4}\right)^{2/3}.$$

Proof. With $K(u) = K_0$, equation (A.2) becomes

$$m_g = \frac{2\beta_g K_0}{\gamma} \int_0^{\Delta_g} \Big(\Delta_g^2 - u^2\Big)\,\mathrm{d}u = \frac{4\beta_g K_0}{3\gamma}\Delta_g^3,$$

which implies the formula for Δ_g.

Similarly, (A.3) becomes

$$T_g - \tau_0 = \frac{\beta_g^2 K_0}{m_g \gamma} \int_0^{\Delta_g} \Big(\Delta_g^4 - u^4\Big)\,\mathrm{d}u = \frac{4\beta_g^2 K_0}{5 m_g \gamma}\Delta_g^5.$$

Substituting the expression for Δ_g yields the closed-form transport wedge. □

Remark A.2 (Comparative statics). The route width satisfies

$$\Delta_g \uparrow \text{ in } m_g, \gamma, \qquad \Delta_g \downarrow \text{ in } K_0, \beta_g.$$

Transport costs satisfy

$$T_g \uparrow \text{ in } m_g, \gamma, \beta_g, \qquad T_g \downarrow \text{ in } K_0.$$

Moreover, the variable part of the transport wedge is homogeneous of degree $-2/3$ in capacity, so

$$-\frac{\partial T_g}{\partial \ln K_0} = \frac{2}{3}(T_g - \tau_0). \tag{A.4}$$

This identity is exact in the constant-capacity benchmark and is useful in the policy section.

A.5 DFS with endogenous transport wedges

Home market. The Home market compares domestic production to imports:

$$w_H a_H(g) \quad \text{vs.} \quad w_F a_F(g) e^{T_{FH}(g)}.$$

The Home import cutoff $\overline{g}$ solves

$$A(\overline{g}) + T_{FH}(\overline{g}) = \varpi. \tag{A.5}$$

Foreign market. The Foreign market compares domestic production to imports from Home:

$$w_F a_F(g) \quad \text{vs.} \quad w_H a_H(g) e^{T_{HF}(g)}.$$

The Home export cutoff $\underline{g}$ solves

$$A(\underline{g}) - T_{HF}(\underline{g}) = \varpi. \tag{A.6}$$

Thus Home exports $g \le \underline{g}$, self-produces $g \in (\underline{g}, \overline{g})$, and imports $g \ge \overline{g}$.

Assumption A.3 (Single crossing with endogenous transport)**.**

$$A'(g) - T'_{HF}(g) < 0, \qquad A'(g) + T'_{FH}(g) < 0 \qquad \forall g \in [0,1].$$

Proposition A.3 (Goods cutoffs)**.** *Under Assumptions A.1 and A.3, there exist unique cutoffs $\underline{g}$ and $\overline{g}$ solving* (A.6) *and* (A.5).

Proof. The export schedule $A(g) - T_{HF}(g)$ and import schedule $A(g) + T_{FH}(g)$ are strictly decreasing by Assumption A.3, so each equation admits at most one solution. Existence follows under the natural boundary conditions that the schedule crosses ϖ on $[0,1]$. □

Comparative statics. For any transport parameter θ,

$$\frac{\mathrm{d}\overline{g}}{\mathrm{d}\theta} = -\frac{\partial_\theta T_{FH}(\overline{g})}{A'(\overline{g}) + T'_{FH}(\overline{g})}, \qquad \frac{\mathrm{d}\underline{g}}{\mathrm{d}\theta} = \frac{\partial_\theta T_{HF}(\underline{g})}{A'(\underline{g}) - T'_{HF}(\underline{g})}.$$

Hence higher transport costs shrink trade and widen the nontraded band.

A.6 Stochastic capacity

Stochastic availability. Let realized capacity be

$$K(x,\omega) = \bar{K}(x)\,\xi(x,\omega), \qquad \xi(x,\omega) \in (0,1].$$

The government chooses baseline capacity $\bar{K}(x)$, then the state ω realizes, and routing is solved in that state. In the common-shock benchmark used below, $\xi(x,\omega) \equiv \xi(\omega)$ is common across route locations. To preserve the closed-form state-by-state solution it is convenient to assume

$$\xi(\omega) \ge \underline{\xi} > 0 \qquad \text{almost surely,}$$

and

$$m_g < \frac{K_0 \underline{\xi}\, \beta_g}{6\gamma},$$

so that the route-support set remains interior in every state.

State-contingent route problem. For imported good g,

$$m_g T_g(\omega) = \min_{f_g(\cdot,\omega)\geq 0} \left\{ \int_0^1 \left[\tau_0 + \beta_g \Big(x - \tfrac{1}{2}\Big)^2\right] f_g(x,\omega)\,\mathrm{d}x + \frac{\gamma}{2}\int_0^1 \frac{f_g(x,\omega)^2}{K(x,\omega)}\,\mathrm{d}x \right\}$$

subject to $\int_0^1 f_g(x,\omega)\,\mathrm{d}x = m_g$. The solution is

$$f_g^\star(x,\omega) = \frac{K(x,\omega)}{\gamma}\left[\mu_g(\omega) - \tau_0 - \beta_g\Big(x - \tfrac{1}{2}\Big)^2\right]_+. \tag{A.7}$$

Common multiplicative availability. The analytically tractable benchmark is

$$K(x,\omega) = K_0\,\xi(\omega).$$

Proposition A.4 (State-contingent route width and transport cost)**.** *Under common stochastic availability,*

$$\Delta_g(\omega) = \left(\frac{3\gamma m_g}{4K_0\xi(\omega)\beta_g}\right)^{1/3},$$

and

$$T_g(\omega) = \tau_0 + C_0\left(\frac{\gamma}{K_0\xi(\omega)}\right)^{2/3}\beta_g^{1/3} m_g^{2/3}.$$

Expected transport cost. Assume

$$\mathbb{E}[\xi^{-2/3}] < \infty.$$

Then taking expectations gives

$$\mathbb{E}[T_g] = \tau_0 + C_0\left(\frac{\gamma}{K_0}\right)^{2/3}\beta_g^{1/3} m_g^{2/3}\,\mathbb{E}[\xi^{-2/3}]. \tag{A.8}$$

Using (A.4) state by state also implies

$$-\frac{\partial\,\mathbb{E}[T_g]}{\partial \ln K_0} = \frac{2}{3}\big(\mathbb{E}[T_g] - \tau_0\big). \tag{A.9}$$

Corollary A.1 (Reliability premium)**.** *Because $z \mapsto z^{-2/3}$ is convex on $(0,\infty)$,*

$$\mathbb{E}[\xi^{-2/3}] > (\mathbb{E}[\xi])^{-2/3}$$

whenever ξ is nondegenerate. Thus stochastic capacity raises expected transport costs even holding mean capacity fixed.

Expected route support. If $\mathbb{E}[\xi^{-1/3}] < \infty$, then

$$\mathbb{E}[\Delta_g] = \left(\frac{3\gamma m_g}{4K_0\beta_g}\right)^{1/3} \mathbb{E}[\xi^{-1/3}],$$

so stochastic capacity widens the expected route area used.

Corollary A.2 (Small-risk approximation)**.** *Let $\bar{\xi} := \mathbb{E}[\xi]$, and suppose ξ fluctuates weakly around $\bar{\xi}$. Then*

$$\mathbb{E}[\xi^{-2/3}] \approx \bar{\xi}^{-2/3}\left(1 + \frac{5}{9}\,\mathrm{CV}(\xi)^2\right), \qquad \mathrm{CV}(\xi)^2 := \frac{\mathrm{Var}(\xi)}{\bar{\xi}^2},$$

and

$$\mathbb{E}[\Delta_g] \approx \Delta_g(\bar{\xi})\left(1 + \frac{2}{9}\,\mathrm{CV}(\xi)^2\right).$$

So to second order, capacity volatility acts like an additional surcharge on expected transport costs that is proportional to the squared coefficient of variation of route availability.

A.7 Expected indirect utility and macro elasticities

Indirect utility. State-contingent indirect utility is

$$V_H(\omega) = \ln E_H - \int_0^{\bar{g}(\omega)} b(g)\ln\big(w_H a_H(g)\big)\,\mathrm{d}g - \int_{\bar{g}(\omega)}^1 b(g)\Big[\ln\big(w_F a_F(g)\big) + T_{FH}(g,\omega)\Big]\mathrm{d}g.$$

Envelope with respect to transport. For any parameter θ that affects the economy only through the transport block,

$$\frac{\mathrm{d}\,\mathbb{E}[V_H]}{\mathrm{d}\theta} = -\,\mathbb{E}\left[\int_{\bar{g}(\omega)}^1 b(g)\,\frac{\partial T_{FH}(g,\omega)}{\partial\theta}\,\mathrm{d}g\right]. \tag{A.10}$$

Proof. Differentiate $V_H(\omega)$ with respect to θ. The terms from the moving cutoff $\bar{g}(\omega)$ vanish because at that margin domestic and imported delivered prices are equal. Taking expectations yields (A.10). □

Aggregate import-transport index. Define

$$\alpha_M := \mathbb{E}\left[\int_{\bar{g}(\omega)}^1 b(g)\,\mathrm{d}g\right]$$

and

$$\mathrm{d}\ln\mathcal{T}_H := \frac{1}{\alpha_M}\,\mathbb{E}\left[\int_{\bar{g}(\omega)}^1 b(g)\,\mathrm{d}T_{FH}(g,\omega)\,\mathrm{d}g\right].$$

Macro welfare elasticity. Define

$$\eta_H^{V,T} := -\frac{\partial\,\mathbb{E}[V_H]}{\partial\ln\mathcal{T}_H}.$$

Under the log benchmark,

$$\eta_H^{V,T} = \alpha_M. \tag{A.11}$$

Reliance and bottleneck objects. Define route reliance for good g in state ω by

$$\theta_g(x,\omega) := \frac{f_g^\star(x,\omega)}{m_g}.$$

The route-level envelope derivative implies

$$-\frac{\partial T_g(\omega)}{\partial \ln \bar{K}(x)} = \theta_g(x,\omega)\,\lambda_g(x,\omega),$$

where

$$\lambda_g(x,\omega) := \frac{\gamma f_g^\star(x,\omega)}{2K(x,\omega)}$$

is the local shadow scarcity.

Now define

$$\Lambda_H(x) := \frac{1}{\alpha_M}\,\mathbb{E}\left[\int_{\overline{g}(\omega)}^1 b(g)\,\theta_g(x,\omega)\,\mathrm{d}g\right], \tag{A.12}$$

$$\mathcal{B}_H(x) := \frac{\mathbb{E}\left[\int_{\overline{g}(\omega)}^1 b(g)\,\theta_g(x,\omega)\lambda_g(x,\omega)\,\mathrm{d}g\right]}{\mathbb{E}\left[\int_{\overline{g}(\omega)}^1 b(g)\,\theta_g(x,\omega)\,\mathrm{d}g\right]}. \tag{A.13}$$

Decomposition. By construction,

$$-\frac{\partial \ln \mathcal{T}_H}{\partial \ln \bar{K}(x)} = \Lambda_H(x)\,\mathcal{B}_H(x). \tag{A.14}$$

A.8 Optimal capacity policy

Policy instrument. Let $s(x)$ denote a subsidy targeted at route location x, and define the pass-through elasticity

$$\chi_K(x) := \frac{\partial \ln \bar{K}(x)}{\partial s(x)}.$$

Government problem. Let the fiscal cost of the subsidy profile be

$$\int_0^1 \mathcal{C}_x\left(s(x)\right)\,\mathrm{d}x, \qquad \mathcal{C}_x'(\cdot) > 0.$$

The government solves

$$\max_{s(\cdot)}\ \overline{W}_0 + p^c\ \mathbb{E}[V_H(s)] - \int_0^1 \mathcal{C}_x\left(s(x)\right)\,\mathrm{d}x.$$

Proposition A.5 (Broad policy rule)**.** *The optimal subsidy at location x satisfies*

$$\mathcal{C}_x'\left(s^\star(x)\right) = p^c\,\eta_H^{V,T}\,\Lambda_H(x)\,\mathcal{B}_H(x)\,\chi_K(x).$$

Proof. The first-order condition is

$$\mathcal{C}_x'(s^\star(x)) = p^c\,\frac{\partial\,\mathbb{E}[V_H]}{\partial s(x)}.$$

Using the chain rule,

$$\frac{\partial \mathbb{E}[V_H]}{\partial s(x)} = \eta_H^{V,T} \left(-\frac{\partial \ln \mathcal{T}_H}{\partial \ln \bar{K}(x)} \right) \chi_K(x).$$

Substitute the decomposition (A.14). □

Interpretation. The broad policy rule has four pieces. The crisis probability p^c scales the expected relevance of transport disruptions. The macro elasticity $\eta_H^{V,T}$ measures how sensitive expected indirect utility is to a proportional change in the average import-transport wedge. The reliance term $\Lambda_H(x)$ captures how much imported expenditure physically uses lane x. The bottleneck term $\mathcal{B}_H(x)$ measures how scarce that lane is, on average, when it is used. Finally, $\chi_K(x)$ converts the subsidy into additional capacity. The rule therefore says that optimal subsidies are largest on lanes that are simultaneously socially important, heavily used, and especially scarce in bad states.

Analytically tractable common-shock benchmark. Now specialize to

$$K(x,\omega) = K_0\, \xi(\omega), \qquad u := |x - \tfrac{1}{2}|.$$

Using Proposition A.4,

$$\Delta_g(\omega) = \left(\frac{3\gamma m_g}{4K_0 \xi(\omega) \beta_g} \right)^{1/3}.$$

Then

$$\theta_g(u,\omega) = \frac{3}{4\Delta_g(\omega)^3} \left[\Delta_g(\omega)^2 - u^2 \right]_+, \tag{A.15}$$

$$\lambda_g(u,\omega) = \frac{\beta_g}{2} \left[\Delta_g(\omega)^2 - u^2 \right]_+. \tag{A.16}$$

Therefore

$$\Lambda_H(u) = \frac{1}{\alpha_M} \mathbb{E} \left[\int_{\bar{g}(\omega)}^1 b(g)\, \frac{3}{4\Delta_g(\omega)^3} \left[\Delta_g(\omega)^2 - u^2 \right]_+ \mathrm{d}g \right], \tag{A.17}$$

$$\mathcal{B}_H(u) = \frac{\mathbb{E}\left[\int_{\bar{g}(\omega)}^1 b(g)\, \frac{3\beta_g}{8\Delta_g(\omega)^3} \left[\Delta_g(\omega)^2 - u^2 \right]_+^2 \mathrm{d}g \right]}{\mathbb{E}\left[\int_{\bar{g}(\omega)}^1 b(g)\, \frac{3}{4\Delta_g(\omega)^3} \left[\Delta_g(\omega)^2 - u^2 \right]_+ \mathrm{d}g \right]}. \tag{A.18}$$

Uniform subsidy: exact rule. If the government uses a uniform subsidy that raises common baseline capacity K_0, define

$$\chi_K := \frac{\partial \ln K_0}{\partial s}.$$

Combining the welfare envelope with (A.9) gives the exact scalar first-order condition

$$\mathcal{C}'(s^\star) = p^c\, \chi_K\, \mathbb{E} \left[\int_{\bar{g}(\omega)}^1 b(g)\, \frac{2}{3} \big(T_g(\omega) - \tau_0 \big)\, \mathrm{d}g \right]. \tag{A.19}$$

Thus the marginal value of capacity is exactly two-thirds of the expenditure-weighted variable part of the transport wedge on the imported range.

Uniform subsidy: factorized approximation. Define

$$\Pi := \mathbb{E}[\xi^{-2/3}],$$

and

$$\Omega_H := \frac{2}{3\alpha_M} \int_{\overline{g}}^{1} b(g)\, C_0 \left(\frac{\gamma}{K_0}\right)^{2/3} \beta_g^{1/3} m_g^{2/3}\, \mathrm{d}g.$$

If one evaluates the imported set at a reference cutoff $\overline{g}$ and neglects the state dependence of that cutoff to first order, then (A.19) reduces to the factorized approximation

$$\mathcal{C}'(s^\star) \approx p^c\, \alpha_M\, \Pi\, \Omega_H\, \chi_K. \tag{A.20}$$

This is the scalar rule reported in the summary. It is best interpreted as a transparent approximation to the exact uniform-capacity rule rather than as a separate exact theorem.

B Quantitative Model: Multi-country Trade and Transport Model

This section records the quantitative environment used in the calibration and counterfactual analysis. Relative to the stylized benchmark, the model here is defined on a finite directed graph with many nodes, country-specific labor mobility, fleet-specific shipping inputs, and soft-min route choice over network paths.

Geography, countries, and fleets. The economy is defined on a directed network $G = (\mathcal{N}, \mathcal{E})$. Nodes $i \in \mathcal{N}$ are production and consumption locations. Directed edges $e = (i \to j) \in \mathcal{E}$ represent feasible shipping links. Let

$$N^+(i) \equiv \{k \in \mathcal{N} : (i \to k) \in \mathcal{E}\}, \qquad N^-(i) \equiv \{k \in \mathcal{N} : (k \to i) \in \mathcal{E}\}.$$

There are M countries indexed $m = 1, \ldots, M$. Nodes partition as $\mathcal{N} = \bigsqcup_{m=1}^{M} \mathcal{N}_m$, and $m(i)$ denotes the unique country such that $i \in \mathcal{N}_{m(i)}$. Country m has total labor endowment $\bar{L}_m$, and labor is perfectly mobile within countries but immobile across countries, implying

$$\sum_{i \in \mathcal{N}_m} L_i = \bar{L}_m, \qquad \frac{u_j w_j}{P_j} = W_m \qquad \forall j \in \mathcal{N}_m.$$

There are N_f fleet origins $n \in \mathcal{F} \equiv \{1, \ldots, N_f\}$. Let Ω denote the state of the world. Fleet availability is encoded by indicators $\Omega_{e,n} \in \{0, 1\}$, where $\Omega_{e,n} = 1$ means fleet n can operate on edge e in state Ω. Rental rates are $\rho_n(\Omega)$.

Preferences, production, and delivered prices. At node j, the CES composite is

$$C_j = \left(\sum_{i \in \mathcal{N}} \varphi_{ij}^{1/\sigma} \, q_{ij}^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, \qquad \sigma > 1,$$

with taste weights $\varphi_{ij} > 0$. Amenities are $u_j > 0$. Let P_j be the CES price index associated with delivered prices P_{ij}^{del}. Indirect utility is

$$U_j = \frac{u_j w_j}{P_j}.$$

Within-country mobility implies $U_j = W_{m(j)}$ for all j.

Each node i produces one variety with linear technology

$$Y_i = A_i L_i,$$

so under perfect competition the factory-gate price is

$$p_i = \frac{w_i}{A_i}.$$

If $\tau_{ij}(\Omega) \geq 1$ denotes the bilateral iceberg trade cost from i to j in state Ω, then the delivered price is

$$P_{ij}^{\text{del}}(\Omega) = p_i \, \tau_{ij}(\Omega).$$

Shipping technology, user prices, and generalized edge costs. Fix an edge $e = (i \to j) \in \mathcal{E}$. Let $S_e(\Omega)$ denote shipping services on edge e in state Ω. Operators rent fleet capacity inputs $\{s_{e,n}(\Omega)\}_{n \in \mathcal{F}}$ at rentals $\rho_n(\Omega)$.

Shipping services are produced by a CES aggregator with elasticity $\eta > 1$:

$$S_e(\Omega) = \left(\sum_{n \in \mathcal{F}} \Omega_{e,n}\, \omega_{e,n}^{1/\eta}\, s_{e,n}(\Omega)^{(\eta-1)/\eta} \right)^{\eta/(\eta-1)}, \qquad \omega_{e,n} > 0.$$

Congestion cost is

$$\Psi_e(S_e) = \frac{\zeta_e}{1+\delta_e} S_e^{1+\delta_e}, \qquad \zeta_e \geq 0,\ \delta_e \geq 0.$$

Taking $\{\rho_n(\Omega)\}$ and the per-unit user price $p_e(\Omega)$ as given, a competitive operator chooses $\{s_{e,n}(\Omega)\}$ to maximize

$$p_e(\Omega) S_e(\Omega) - \sum_n \rho_n(\Omega)\, s_{e,n}(\Omega) - \Psi_e(S_e(\Omega)).$$

Define the dual unit cost of the optimized fleet bundle, excluding congestion, by

$$c_e(\Omega) \equiv \left(\sum_{n \in \mathcal{F}} \Omega_{e,n}\, \omega_{e,n}\, \rho_n(\Omega)^{1-\eta} \right)^{1/(1-\eta)}. \tag{B.1}$$

Under perfect competition, the per-unit user price equals marginal cost:

$$p_e(\Omega) = c_e(\Omega) + \zeta_e\, S_e(\Omega)^{\delta_e}. \tag{B.2}$$

Let $\bar{\kappa}_e(\Omega) > 0$ denote the technological component of shipping on edge e, and let $v_e > 0$ be the average value of shipments transported on edge e. Define the generalized iceberg cost on edge e as

$$\kappa_e(\Omega) \equiv \bar{\kappa}_e(\Omega) \left(1 + \frac{p_e(\Omega)}{v_e} \right) = \bar{\kappa}_e(\Omega) \frac{p_e(\Omega) + v_e}{v_e}. \tag{B.3}$$

Define the shipping-price share in the ad-valorem term

$$\psi_e(\Omega) \equiv \frac{p_e(\Omega)}{p_e(\Omega) + v_e} \in (0,1), \tag{B.4}$$

the CES fleet cost share

$$\pi_{e,n}(\Omega) \equiv \frac{\Omega_{e,n}\, \omega_{e,n}\, \rho_n(\Omega)^{1-\eta}}{\sum_{h \in \mathcal{F}} \Omega_{e,h}\, \omega_{e,h}\, \rho_h(\Omega)^{1-\eta}}, \qquad \sum_n \pi_{e,n}(\Omega) = 1, \tag{B.5}$$

and the fleet-bundle share in the user price

$$\vartheta_e(\Omega) \equiv \frac{c_e(\Omega)}{p_e(\Omega)} \in (0,1]. \tag{B.6}$$

Soft-min routing and edge-usage probabilities. Let $\mathcal{R}(i \to j)$ denote the set of all finite directed walks from i to j in G, allowing repeats. For a route $r \in \mathcal{R}(i \to j)$, let $\mathcal{E}(r)$ denote the multiset of edges along r. Define the route iceberg cost as

$$\kappa_r(\Omega) \equiv \prod_{e \in \mathcal{E}(r)} \kappa_e(\Omega).$$

Shippers choose routes under a logit structure with scale parameter $\phi > 1$. The implied inclusive-value

bilateral iceberg cost $\tau_{ij}(\Omega)$ is defined by

$$\tau_{ij}(\Omega)^{1-\phi} = \sum_{r \in \mathcal{R}(i \to j)} \kappa_r(\Omega)^{1-\phi}. \tag{B.7}$$

Proposition B.1 (Soft-min routing recursion). *For all $i, j \in \mathcal{N}$,*

$$\tau_{ij}(\Omega)^{1-\phi} = \mathbf{1}\{i = j\} + \sum_{k \in N^+(i)} \kappa_{ik}(\Omega)^{1-\phi}\, \tau_{kj}(\Omega)^{1-\phi}. \tag{B.8}$$

Proof. Every walk from i to j is either the empty walk, only possible if $i = j$, or begins with a unique first step $(i \to k)$ for some $k \in N^+(i)$. For any fixed k, the remainder of the walk is a walk from k to j, so the contribution of walks beginning with $(i \to k)$ equals $\kappa_{ik}^{1-\phi}\tau_{kj}^{1-\phi}$. Summing over k gives (B.8). □

Proposition B.2 (Edge-usage probability). *Fix origin–destination (i, j) and an edge $(k \to \ell)$. Under the routing structure above, the probability that a shipment from i to j uses edge $(k \to \ell)$ in state Ω is*

$$\theta_{(k\to\ell)}(i \to j; \Omega) = \frac{\tau_{ik}(\Omega)^{1-\phi}\, \kappa_{k\ell}(\Omega)^{1-\phi}\, \tau_{\ell j}(\Omega)^{1-\phi}}{\tau_{ij}(\Omega)^{1-\phi}}. \tag{B.9}$$

Proof. By (B.7), the unnormalized logit weight on a route r from i to j is proportional to $\kappa_r^{1-\phi}$. The total weight of all routes from i to j that use edge $(k \to \ell)$ can be decomposed into a sub-route from i to k, the edge $(k \to \ell)$, and a sub-route from ℓ to j. The sum of weights over such sub-routes factorizes as $\tau_{ik}^{1-\phi}\kappa_{k\ell}^{1-\phi}\tau_{\ell j}^{1-\phi}$. Dividing by the total weight $\tau_{ij}^{1-\phi}$ yields (B.9). □

Pass-through, propagation, and welfare aggregation.

Proposition B.3 (Pass-through from rentals to lane iceberg costs). *Fix a state Ω and an edge e. Consider a small proportional change in a fleet rental $\rho_n(\Omega)$. Holding $\bar{\kappa}_e(\Omega)$, v_e, and $S_e(\Omega)$ fixed, the resulting local pass-through is*

$$\frac{\partial \ln \kappa_e(\Omega)}{\partial \ln \rho_n(\Omega)} = \psi_e(\Omega)\, \vartheta_e(\Omega)\, \pi_{e,n}(\Omega), \tag{B.10}$$

where ψ_e is defined in (B.4), ϑ_e in (B.6), and $\pi_{e,n}$ in (B.5). In the benchmark case without congestion, $\vartheta_e(\Omega) = 1$ and

$$\frac{\partial \ln \kappa_e(\Omega)}{\partial \ln \rho_n(\Omega)} = \psi_e(\Omega)\, \pi_{e,n}(\Omega).$$

Proof. From (B.3), $\ln \kappa_e = \ln \bar{\kappa}_e + \ln(p_e + v_e) - \ln v_e$, so

$$\frac{\partial \ln \kappa_e}{\partial \ln p_e} = \frac{p_e}{p_e + v_e} = \psi_e.$$

From (B.2), holding S_e fixed implies $\partial p_e / \partial c_e = 1$, hence

$$\frac{\partial \ln p_e}{\partial \ln c_e} = \frac{c_e}{p_e} = \vartheta_e.$$

Finally, from the CES dual cost (B.1), the standard share formula gives $\partial \ln c_e / \partial \ln \rho_n = \pi_{e,n}$. Multiplying the three derivatives yields (B.10). □

Proposition B.4 (Bilateral trade-cost linearization). *Fix (i, j) and a state Ω. For small changes $\{d \ln \kappa_e(\Omega)\}_{e \in \mathcal{E}}$, the induced change in the bilateral iceberg cost satisfies*

$$d \ln \tau_{ij}(\Omega) = \sum_{e \in \mathcal{E}} \theta_e(i \to j; \Omega)\, d \ln \kappa_e(\Omega), \tag{B.11}$$

where $\theta_e(i \to j; \Omega)$ is the probability that an $i \to j$ shipment uses edge e.

Proof. From (B.7),

$$\tau_{ij}^{1-\phi} = \sum_{r \in \mathcal{R}(i \to j)} \kappa_r^{1-\phi}.$$

Differentiate and use $\ln \tau_{ij} = \frac{1}{1-\phi} \ln \left(\sum_r \kappa_r^{1-\phi} \right)$:

$$d \ln \tau_{ij} = \sum_r \frac{\kappa_r^{1-\phi}}{\sum_{r'} \kappa_{r'}^{1-\phi}}\, d \ln \kappa_r.$$

Because $d \ln \kappa_r = \sum_{e \in \mathcal{E}} n_{r,e}\, d \ln \kappa_e$, where $n_{r,e}$ is the number of times route r uses edge e,

$$d \ln \tau_{ij} = \sum_{e \in \mathcal{E}} \left(\sum_r \text{route weight} \cdot n_{r,e} \right) d \ln \kappa_e.$$

The inner term is the expected usage of edge e under the route-choice distribution, which equals $\theta_e(i \to j)$. □

Remark B.1 (General welfare statement). With indirect utility $U_j = u_j w_j / P_j$ and within-country equalization $U_j = W_{m(j)}$, small welfare changes can always be written as combinations of changes in wages, amenities, and the price index. The paper's empirical aggregation is a benchmark approximation that focuses on delivered import prices and scales by imports as a share of GDP.

Proposition B.5 (Import-share welfare aggregation (benchmark)). *Let Y denote GDP and let $s_{\text{imp}}^{\text{GDP}}$ denote the share of imports in GDP. Let w_e denote the share of aggregate imports that use lane e, with $\sum_e w_e = 1$. If a policy induces small proportional lane iceberg changes $\{\Delta \ln \kappa_e\}_e$ and these changes pass through to delivered import prices one-for-one in the benchmark aggregation, then equivalent variation satisfies*

$$\frac{\Delta EV}{Y} \approx - s_{\text{imp}}^{\text{GDP}} \sum_{e \in \mathcal{E}} w_e\, \Delta \ln \kappa_e. \tag{B.12}$$

Proof. In the benchmark, the proportional change in the aggregate delivered import price index is approximated by the import-share-weighted average of lane price changes, $\sum_e w_e \Delta \ln \kappa_e$. The change in the cost of purchasing the baseline import bundle is then approximately $-M_{\text{imp}} \sum_e w_e \Delta \ln \kappa_e$, where $M_{\text{imp}} = s_{\text{imp}}^{\text{GDP}} Y$. Dividing by Y yields (B.12). □

Risk adjustment and equilibrium.

Proposition B.6 (Certainty-equivalent generalized edge cost). *Suppose $\ln \kappa_e$ is normally distributed with mean μ_e^κ and variance s_e^2. Under CARA preferences with coefficient $\chi > 0$, the certainty-equivalent log cost is*

$$\ln \kappa_e^{CE} = \mu_e^\kappa + \frac{\chi}{2} s_e^2, \qquad \kappa_e^{CE} = \exp\left(\mu_e^\kappa + \frac{\chi}{2} s_e^2 \right). \tag{B.13}$$

Proof. Standard CARA-normal certainty-equivalent logic: the certainty equivalent equals the mean plus $\chi/2$ times the variance in the relevant log cost. □

Remark B.2 (Routing under risk). All routing expressions above apply with $\kappa_e(\Omega)$ replaced by the certainty-equivalent object $\kappa_e^{CE}(\Omega)$.

A state-contingent competitive equilibrium consists of prices and allocations such that: (i) households maximize utility and, within each country, relocate to equalize $u_j w_j / P_j$ subject to $\sum_{i \in \mathcal{N}_m} L_i = \bar{L}_m$; (ii) firms choose labor, produce $Y_i = A_i L_i$, and set $p_i = w_i / A_i$; (iii) shipping operators choose $\{s_{e,n}(\Omega)\}$ on each edge to maximize profits, yielding (B.1)–(B.2); (iv) routing probabilities and $\tau_{ij}(\Omega)$ satisfy (B.8) and (B.9); (v) goods markets clear; (vi) shipping-services markets clear; (vii) fleet-capacity markets clear; and (viii) labor markets clear.

C Quantitative Model: Hat Algebra and Computation

This section records the exact-hat system used in the quantitative counterfactuals and the numerical algorithm used to solve it. The key simplification is that the counterfactual equilibrium can be written in terms of directed edge traffic and baseline fleet shares, so the implementation does not require recomputing all bilateral route objects in levels.

Hats, unknowns, and normalization. Using the notation above, define hats by $\hat{x} \equiv x/x^0$, where x^0 denotes the baseline. The counterfactual equilibrium is solved in the core unknowns

$$\{\hat{w}_i\}_{i\in\mathcal{N}}, \qquad \{\hat{y}_i\}_{i\in\mathcal{N}}, \qquad \{\hat{W}_m\}_{m=1}^{M}, \qquad \{\hat{\kappa}_e\}_{e\in\mathcal{E}}, \qquad \{\hat{\rho}_n\}_{n=1}^{N_f},$$

where $\hat{w}_i$ are wages, $\hat{y}_i$ are node nominal-income hats, $\hat{W}_m$ are country welfare multipliers, $\hat{\kappa}_e$ are generalized edge iceberg costs, and $\hat{\rho}_n$ are fleet rental rates.

The outward income recursion is homogeneous, so $\hat{y}$ is pinned down only up to scale. Let $Y_i^{out,0} \equiv w_i^0 L_i^0$ and define income weights

$$\alpha_i \equiv \frac{Y_i^{out,0}}{\sum_{j\in\mathcal{N}} Y_j^{out,0}}.$$

We impose the normalization

$$\sum_{i\in\mathcal{N}} \alpha_i \, \hat{y}_i = 1. \tag{C.1}$$

Within each country m, mobility implies $u_i w_i / P_i = W_m$ for all $i \in \mathcal{N}_m$, hence in hats

$$\hat{P}_i = \frac{\hat{u}_i \, \hat{w}_i}{\hat{W}_{m(i)}}.$$

Define node nominal income in levels as $Y_i^{out} \equiv w_i L_i$ and its hat as

$$\hat{y}_i \equiv \widehat{Y}_i^{out} = \hat{w}_i \, \hat{L}_i, \qquad \text{so that} \qquad \hat{L}_i = \frac{\hat{y}_i}{\hat{w}_i}. \tag{C.2}$$

Country labor endowments are fixed:

$$\sum_{i\in\mathcal{N}_m} L_i^0 \, \frac{\hat{y}_i}{\hat{w}_i} = \sum_{i\in\mathcal{N}_m} L_i^0 \qquad \forall m. \tag{C.3}$$

Traffic-based outer block. Let Ξ_{ik} denote value traffic, that is, expenditure on shipping services on edge $(i \to k)$. The baseline equilibrium delivers convex weights

$$\omega_i^{out,loc}, \quad \{\omega_{ik}^{out,edge}\}_{k\in N^+(i)}, \qquad \omega_i^{in,loc}, \quad \{\omega_{ki}^{in,edge}\}_{k\in N^-(i)},$$

satisfying

$$\omega_i^{out,loc} + \sum_{k\in N^+(i)} \omega_{ik}^{out,edge} = 1, \qquad \omega_i^{in,loc} + \sum_{k\in N^-(i)} \omega_{ki}^{in,edge} = 1.$$

To eliminate bilateral trade objects and express equilibrium purely in terms of edge traffic, adopt the aligned case in which the curvature governing routing and CES demand coincide, so the relevant power is $1-\sigma$ throughout.

Then the outgoing edge-traffic hat can be written as

$$\widehat{\Xi}_{(i\to k)}(\hat{w},\hat{y},\hat{\kappa}) \equiv \hat{\kappa}_{ik}^{1-\sigma}\,\hat{A}_i^{\sigma-1}\hat{w}_i^{1-\sigma}\,\hat{A}_k^{1-\sigma}\hat{w}_k^{\sigma-1}\hat{y}_k. \tag{C.4}$$

For each node i, the outward equation is

$$\hat{y}_i = \omega_i^{out,loc}\Big[\hat{W}_{m(i)}^{1-\sigma}\,\hat{A}_i^{\sigma-1}\,\hat{u}_i^{\sigma-1}\,\hat{y}_i\Big] + \sum_{k\in N^+(i)} \omega_{ik}^{out,edge}\,\widehat{\Xi}_{(i\to k)}(\hat{w},\hat{y},\hat{\kappa}), \tag{C.5}$$

and the inward equation is

$$1 = \omega_i^{in,loc}\Big[\hat{W}_{m(i)}^{1-\sigma}\,\hat{A}_i^{\sigma-1}\,\hat{u}_i^{\sigma-1}\Big] + \sum_{k\in N^-(i)} \omega_{ki}^{in,edge}\Big[\hat{\kappa}_{ki}^{1-\sigma}\,\hat{u}_i^{\sigma-1}\hat{w}_i^{\sigma-1}\,\hat{u}_k^{1-\sigma}\hat{w}_k^{1-\sigma}\Big]. \tag{C.6}$$

Transport block and counterfactual equilibrium. Let v_e denote value per service so that $\Xi_e = v_e S_e$. Then

$$\hat{S}_e = \frac{\widehat{\Xi}_e(\hat{w},\hat{y},\hat{\kappa})}{\hat{v}_e}. \tag{C.7}$$

Let $\pi_{e,n}$ be the baseline cost share of fleet n on edge e in the CES aggregator, and let $\widehat{\Omega}_{e,n}$ be a counterfactual shifter that scales the effective unit cost of using fleet n on edge e. Define the fleet-bundle hat

$$\widehat{\kappa_e^{fleet}}(\hat{\rho}) \equiv \left(\sum_{n=1}^{N_f} \pi_{e,n}\,(\widehat{\Omega}_{e,n}\hat{\rho}_n)^{\,1-\eta}\right)^{\frac{1}{1-\eta}}. \tag{C.8}$$

Let $\psi_e \in [0,1]$ denote the baseline share of fleet costs in the user price, and let $\xi_e \in [0,1)$ denote the baseline share of the ad-valorem term in generalized costs. Then the generalized edge-cost hat satisfies

$$\hat{\kappa}_e = \widehat{\bar{\kappa}}_e\left[(1-\xi_e) + \xi_e\left(\psi_e\,\widehat{\kappa_e^{fleet}}(\hat{\rho}) + (1-\psi_e)\,\hat{S}_e^{\delta_e}\right)\right], \qquad \forall e\in\mathcal{E}, \tag{C.9}$$

where $\widehat{\bar{\kappa}}_e$ is an edge technological shifter and $\delta_e \geq 0$ governs congestion curvature.

Let $s_{n,e}$ denote the baseline usage share of fleet n on edge e, and let $\hat{\mu}_n$ be a fleet-specific supply shifter. Fleet rentals satisfy

$$\hat{\rho}_n^{1+\eta} = \hat{\mu}_n \sum_{e\in\mathcal{E}} s_{n,e}\,\hat{S}_e\,\widehat{\Omega}_{e,n}\left(\widehat{\kappa_e^{fleet}}(\hat{\rho})\right)^{\eta}, \qquad \forall n\in\{1,\dots,N_f\}. \tag{C.10}$$

A counterfactual equilibrium is a collection of hats satisfying: (i) country labor constraints (C.3); (ii) nominal normalization (C.1); (iii) the traffic-based outer equations (C.5)–(C.6); and (iv) the transport block (C.7)–(C.10) with traffic defined by (C.4).

Implementing targeted shocks. The quantitative exercises map targeted shocks into $(\widehat{\Omega}_{e,n}, \hat{\mu}_n, \widehat{\bar{\kappa}}_e)$.

- **Ad valorem taxes or surcharges on fleet n.** A tax T_n on fleet n that applies on a set of edges $\mathcal{E}^{tax}$ can be represented by

$$\widehat{\Omega}_{e,n} = 1 + T_n \quad \text{for } e\in\mathcal{E}^{tax}, \qquad \widehat{\Omega}_{e,n} = 1 \quad \text{otherwise.}$$

- **Disruptions to availability.** A disruption that makes fleet n unavailable on edges $\mathcal{E}^{ban}$ is implemented by setting $\widehat{\Omega}_{e,n}$ to a very large number, or equivalently by dropping the term from the CES sum, for $e \in \mathcal{E}^{ban}$.
- **Fleet supply shifts.** A change in the convex supply slope for fleet n is represented by $\hat{\mu}_n \neq 1$ in (C.10).
- **Edge technological shifts.** Route-specific operational changes, such as port disruptions, are represented by $\widehat{\bar{\kappa}}_e \neq 1$ in (C.9).

Numerical solution. A stable implementation mirrors the model's economic nesting: the outer loop updates $\hat{W}$ to satisfy country labor endowments; the middle loop solves $(\hat{w}, \hat{y})$ given $(\hat{W}, \hat{\kappa})$; and the inner loop solves $(\hat{\rho}, \hat{\kappa})$ given $(\hat{w}, \hat{y})$.

Step 0 (initialization). Initialize $\hat{W}_m = 1$ for all m, $\hat{\kappa}_e = 1$ for all e, and positive vectors $\hat{w}, \hat{y}$ with $\sum_i \alpha_i \hat{y}_i = 1$.

Step 1 (middle loop: real side given $\hat{W}, \hat{\kappa}$). Given $(\hat{W}, \hat{\kappa})$, update wages $\hat{w}$ by iterating on (C.6) using damping or Anderson acceleration in $\ln \hat{w}$. Then, holding $(\hat{W}, \hat{w}, \hat{\kappa})$ fixed, update incomes $\hat{y}$ by iterating on the linear outward recursion (C.5). After each iteration, reimpose the nominal normalization

$$\hat{y} \leftarrow \frac{\hat{y}}{\sum_i \alpha_i \hat{y}_i}.$$

Finally set $\hat{L}_i = \hat{y}_i / \hat{w}_i$.

Step 2 (inner loop: transport given $\hat{w}, \hat{y}, \hat{\kappa}$). Compute edge traffic $\widehat{\Xi}_e$ using (C.4) and services $\hat{S}_e$ using (C.7). Solve fleet rentals $\hat{\rho}$ from (C.10), using a fixed point or Anderson acceleration, then update edge costs $\hat{\kappa}$ from (C.9). Iterate Steps 1 and 2 until $\hat{\kappa}$ is a fixed point.

Step 3 (outer update: welfare multipliers). Compute implied total labor in each country

$$\widehat{L}_m^{tot} \equiv \frac{\sum_{i \in \mathcal{N}_m} L_i^0 (\hat{y}_i / \hat{w}_i)}{\sum_{i \in \mathcal{N}_m} L_i^0},$$

and update welfare multipliers by a damped tatonnement:

$$\ln \hat{W}_m \leftarrow \ln \hat{W}_m + \gamma_W \ \ln(\widehat{L}_m^{tot})^{-1}.$$

Convergence requires $\widehat{L}_m^{tot} \approx 1$ for all m and small residuals in the outer and transport blocks.

With within-country mobility, $\hat{W}_m$ is the country welfare change associated with the counterfactual. For the benchmark aggregation used in the reduced-form surcharge exercise, the same logic yields the import-share approximation reported in the main text.