2158

Discussion Papers

Quantitative Trade with Ships

Sonali Chowdhry, Inga Heiland, Hendrik Mahlkow

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

Quantitative trade with ships [*]

Sonali Chowdhry, Inga Heiland, Hendrik Mahlkow

MARCH 2026

This paper highlights an underexplored margin of heterogeneity that shapes resilience to disruptions in global maritime trade - the differential reliance of countries and sectors on specific categories of vessels. We combine US bills of lading records with ship registry and AIS-based port call data to document new stylized facts on vessel deployment, including switching patterns across ships, country specialization in shipbuilding, and the composition of fleets serving different country pairs. Exploiting the 2016 Panama Canal expansion as a quasi-natural experiment, we further provide the first direct estimate for the elasticity of substitution between vessels across size classes. Building on the empirical evidence, we then introduce endogenous vessel choice into a quantitative general equilibrium trade model that features multiple transport modes and a global market for shipping services. The model allows us to quantify the trade and welfare effects of two recent policy proposals that target specific ship types, namely, fees for Chinese-built vessels entering US ports and the inclusion of the maritime transport sector in the EU Emission Trading System.

 Keywords: Maritime transport, Quantitative general equilibrium trade models, EU ETS, Port fees, China

JEL-Classification: F13, F14, F52, R41

1. Introduction

Maritime transport is the backbone of international trade. Yet heightened geopolitical tensions and extreme weather events have renewed concerns about the vulnerability and resilience of the shipping network. A defining feature of the major shocks and policy changes facing global maritime shipping is that they affect particular segments of the merchant fleet disproportionately. For instance, a new port fees proposed by the US in 2025 amid escalating US-China trade tensions exclusively targets Chinese-built vessels serving importers and exporters in the US. Climate policies like the EU Emission Trading System (EU ETS) hit first and foremost the emissions-intensive segment of the global fleet serving European countries. Likewise, drought-induced restrictions in the Panama Canal in 2023 were binding only for larger vessels and the COVID-19 induced collapse in port handling capacity disrupted services for high-capacity container vessels.

Hence, the resilience of global shipping and the economic consequences of shocks and policies targeting maritime transport depend crucially on their incidence across ship types and on the flexibility of merchant fleets. Despite this, we know surprisingly little about the fleets serving particular countries and routes. Similarly, little is known about carriers' flexibility with regard to redeploying vessels across routes, and the substitutability of vessels with different characteristics.

In this paper, we open up this largely unexamined margin of vessel heterogeneity in global maritime transport. We study the endogenous selection of ship types across routes as a new margin of adjustment to vessel-type specific policies and provide novel descriptive and empirical evidence for the importance of this margin. Moreover, we develop a quantitative framework that allows us to flexibly quantify the general equilibrium trade and welfare effects of vessel-type specific policies and shocks.

To do so, we first construct a new database that combines US bill of lading records, global ship registries and AIS-based port call information. With this data, we can observe how the composition of vessels serving the US evolves across firms, products, source countries, and carriers. Such an analysis has not been feasible with conventional trade data thus far, which lack details either on the transport mode, vessel identifiers or the detailed mapping between vessels and goods.

Our first contribution with this combined data is to uncover a series of new stylized facts on how maritime transportation adjusts at the vessel level. Here, we show that port and carrier choices are highly persistent for a given importer–product–origin country triplet. By contrast, the vessel margin is fluid with ships being rotated frequently across consecutive shipments. Yet this flexibility is tightly circumscribed. Vessel switching is overwhelmingly within the same ship type and size class, indicating that substitution is feasible only within narrow technological segments of the fleet.

These segments are themselves shaped by long-run shipbuilding specialization. Japan, South Korea and European shipyards occupy different market positions within the ship type and size space, with China showing dramatic growth in newly constructed ships across all segments in the past three decades. Finally, we show that these China-built vessels are deeply embedded in global trade. They serve a wide range of country pairs, occupy substantial shares in the fleets of all major carriers and serve a quarter of all US importers.

In addition to these stylized facts, our second key contribution is to estimate the elasticity of substitution across vessels. We do so by exploiting the 2016 expansion of the Panama Canal as a quasi-natural experiment. This expansion added a new set of locks with strict dimensional thresholds, which generates sharp treatment variation across vessels. With an event study specification, our results show rapid upgrading of vessel sizes and reallocation of US imports from Panamax to larger Post-Panamax ships on treated routes in addition to a decline in maritime freight charges. Taking the ratio of these value and price responses delivers the first direct estimate of the elasticity of substitution across vessel size classes in global shipping.

Our third contribution is to develop a quantitative trade model with multiple transport modes, endogenous choice of vessel types, and a global market for vessel services. Our model provides a flexible and tractable framework to analyze the general equilibrium consequences of a range of policy measures and shocks that differentially affect vessels of different types. We employ the model to study two pertinent policies: (i) fees for Chinese-built vessels entering US ports as proposed by the Trump administration and (ii) the extension of the EU Emission Trading System (EU ETS), requiring vessel operators to purchase allowances for emissions related to voyages to or from ports in the European Economic Area (EEA) as well as voyages between EEA ports.

Crucially, the effects of such policies extend far beyond the US and the EEA, respectively. By changing the relative demand for ships built in China in the first scenario and the relative demand for emission-intensive vessels in the second, these policies generate global spillovers through changes in the price of vessel services. Our model is designed to capture the direct and indirect effects of such policies targeting maritime shipping, allowing for heterogeneity in the incidence of the policy measure, global spillovers through the shipping market as well as more standard general equilibrium effects working through the goods markets.

To preview the results, we find that the economic costs of US port fees on Chinese built vessels are borne largely by the US economy, through increased prices of imported final goods for consumers and imported intermediates for producers. Price increases are concentrated among goods that are typically transported with bulk ships, for which China-built vessels dominate the market. The port fees also generate sizable positive effects for other countries, in particular countries that use Chinese-built vessels intensively for their global imports and exports.

Our simulation of the inclusion of maritime transport in the European ETS, which imposes heterogeneous costs based on vessel emission intensity, yields strong negative economic effects for the EEA. EEA members face higher prices of imports and exports, leading to a loss of competitiveness relative to other countries. The trade effects are strongest for those EEA countries whose imports and exports rely on small emission-intensive vessels. Furthermore, the policy generates important spillover effects as global prices of services of emission-intensive vessels drop compared to less emission-intensive vessels and compared to other transport modes. This results in the redeployment of the small, emission-intensive vessels on non-EEA trade routes.

Our scenarios highlight how unilateral transport policies have global redistributive effects through endogenous adjustments in the shipping market. While we focus on builder country and emission intensity of vessels in this paper, our model readily extends to other observable ship characteristics such as its flag state or age. As such, the framework is well-suited to analyze the growing array of vessel-specific shocks currently shaping international trade, ranging from geopolitical sanctions to the scrapping of old vessels.

Contributions to related literature

This paper contributes to the literature on international trade and transportation in four crucial dimensions. First, it sheds light on a new and more granular margin of adjustment in the transport sector by examining the deployment of heterogeneous vessels across import sectors and maritime routes. In doing so, we extend prior work which has investigated alternative margins such as the selection of trade flows into air, sea or land-based transportation (Hummels and Schaur 2013; Harrigan 2010), optimal maritime shipping routes (Ganapati, Wong, and Ziv 2024; Heiland et al. 2025; Do et al. 2025), choice across carriers (Cristoforoni et al. 2025; Ardelean and Lugovskyy 2023) and ports (Koenig et al. 2024; Verschuur, Koks, and Hall 2022). However, the empirical evidence has been scarce on reallocation within maritime shipping along the vessels margin. We address this gap by analyzing substitutions across ship categories, size classes and builder countries using transaction-level data that reports vessel identifiers.

Second, this paper exploits the expansion of the Panama Canal in 2016 to provide a novel estimate of the elasticity of substitution between vessels across size classes. This complements Tolva (2024) who estimates substitution elasticities but across modes of transportation using a Wald difference-in-differences estimation strategy (De Chaisemartin and d'Haultfoeuille 2018). Furthermore, by showing how the Panama Canal expansion altered the composition of fleets serving US imports, we also add to ongoing research on the economic effects of port and canal infrastructure investments (Bonadio 2024; Brancaccio, Kalouptsidi, and Papageorgiou 2024; Feyrer 2021).

Third, we add to the literature studying the endogenous determination of transportation cost in quantitative spatial models (Allen and Arkolakis 2022; Brancaccio, Kalouptsidi, and Papageorgiou

2020; Asturias 2020; Wong 2022; Ganapati, Wong, and Ziv 2024). We build a nested choice over transportation modes following Tolva (2024) into the Caliendo and Parro (2015) framework, add a nested choice over different types of vessels, and model explicitly the global and interdependent markets for the services of different types of vessels. Relative to Ludwig (2025), who models the allocation of heterogeneous container ships to circular routes and studies how unilateral carbon policy reshapes network efficiency and emissions, our framework differs in two key ways. First, we cover multiple vessel types (containers, tankers, bulk) and transport modes rather than focusing on containers alone. Second, we introduce a builder-country dimension and an endogenous global market for vessel services, which generates policy spillovers through ship-service prices and ship-service income rather than through network reallocation.

Our results on the US proposal to impose port call fees on China-built vessels also expands the scope of recent work on the geoeconomic fragmentation of global trade and the weaponisation of economic interdependence (Clayton et al. 2025; Gopinath et al. 2025; Farrell and Newman 2019). This rapidly growing literature has thus far focused on the use of punitive measures that target goods and services directly, such as import sanctions (Aytun, Hinz, and Özgüzel 2025; Chowdhry et al. 2024; Felbermayr et al. 2020), export controls (Alfaro et al. 2025) and banking restrictions (Drott, Goldbach, and Nitsch 2024). Here, we highlight the impact of policies that intervene in maritime transportation and potentially segment the global shipping market along geopolitical dimensions. Such interventions operate through different channels than trade sanctions, as they directly alter global transportation costs for countries and sectors.

The rest of the paper is structured as follows. Section 2 introduces the main datasets used in the paper. Section 3 leverages these datasets to document a series of novel stylized facts on vessel turnover, shipbuilding and fleet deployment across country pairs. In Section 4, we report our reduced-form findings on the elasticity of substitution across vessels alongside multiple robustness checks. Section 5 embeds these empirical findings into a new quantitative trade model that features choice across modes of transportation as well as vessel types. Section 6 leverages this model to simulate counterfactuals on the proposed US port fees on China-built vessels and carbon pricing for polluting vessels while Section 7 concludes with an outlook for future research.

2. Data

Conventional trade databases are inadequate for analyzing substitution across vessel types within the maritime sector. Standard datasets such as BACI, widely used in gravity estimations, provide detailed trade flows by country pair and product but do not distinguish between modes of transport. UN Comtrade and US Census data offers such a disaggregation across modes, but lacks information on the types of vessels employed conditional on the use of maritime shipping.

Finally, AIS data contain rich information on vessel movements but do not identify the specific goods transported or their detailed origins beyond the port of departure. To quantify the exposure of trade flows to different vessel types, we therefore construct a rich database that links firms, vessels, products, and source countries. This dataset is built by merging three distinct data sources, described in detail below.

Bill of lading data: The first database consists of bill of lading records that capture transaction-level information on US maritime imports. Each record contains a unique identifier for the US consignee or importing entity, the set of HS 6-digit products shipped as well as the origin country, value and volume of the transaction. Crucially, each transaction reports the International Maritime Organization (IMO) identifier of the vessel used for transportation, the foreign port of departure and US port of arrival. In terms of coverage, the data spans more than 140 million transactions recorded between January 2015 and December 2024. As an example, data from the latest year (2024) spans more than 600,000 US importing entities, sourcing over 5700 HS-6 digit products via 6000 different vessels arriving across multiple US ports.

Vessel characteristics: The second database provides detailed information on vessel characteristics and sheds light on the composition of the global shipping fleet. Each vessel, identified by its IMO number, is linked to its country and year of build, as well as several design and size attributes which determine its carrying capacity and fuel efficiency. In total, the database covers more than 67,000 vessels, including container ships, dry bulkers, tankers and other vessel types.

Vessel port calls: The third database tracks the geographical positions and imputed port calls of ships over time. Merging this data with vessel characteristics (using IMO numbers) allows us to examine the deployment of different vessel types across trade routes. Following unilateral policy changes affecting vessel choice, this data also allows us to examine how bilateral iceberg trade costs between third countries may change depending on their exposure to targeted vessel types. Here, we use data for the latest full-year available (2024) which shows port calls made by nearly 10,000 different vessels across 3600 ports located in 180 countries.

3. Stylised Facts

This unique combination of bill of lading data, vessel characteristics, and port call information uncovers a previously overlooked margin of adjustment in the international trade and transportation literature. It enables a detailed examination of how flexibly maritime transport responds to shocks not only through changes in shipping routes, but also through adjustments in the composition of vessels deployed across those routes. We begin first by presenting a series of novel stylized facts on the use of vessels types across US importing firms and country pairs.

Table 1 — Probability of changing margins across US shipments

Vessel	Builder	Carrier	Port	Probability
		No switch		
–	–	–	–	0.282
		Single-margin switches		
✓	–	–	–	0.232
–	–	✓	–	0.019
–	–	–	✓	0.001
		Multiple-margin switches		
✓	✓	–	–	0.223
✓	✓	✓	–	0.109
✓	–	✓	–	0.070
✓	✓	✓	✓	0.022
✓	✓	–	✓	0.020
✓	–	–	✓	0.012
✓	–	✓	✓	0.011
0.699	0.374	0.231	0.066	1.000

Note: This table reports the probability of different types of switches over consecutive shipments of the same US importer–product–origin country triplet. Calculations are based on US maritime imports spanning 2020-2024.

Stylized fact 1. **Port and carrier choices are sticky but vessels change frequently across firm-level shipments**

Using detailed US bill of lading data, we can track how often individual importing firms experience switches between vessels, carriers and ports of unlading, even when sourcing the same product from the same origin country. Drawing on shipment-level observations from the latest five years (2020-2024), we compute the probability of such switching across consecutive transactions by the same firm–product–origin country combination. We further distinguish between several categories of switching behavior. These include instances where the firm sees a change of vessel, builder country of the vessel, carrier company and/or US port of unlading. Table 1 reports the distribution of switching probabilities depending on how many and which margins change.[1]

Vessel changes emerge as the dominant margin of adjustment. In nearly 70% of cases, the vessel used differs from the previous shipment when sourcing the same product from the same origin country. Moreover, only 30% of these vessel changes involve switching to a different carrier. The majority of vessel turnover is thus driven by *within-carrier* redeployment of fleets.[2] Thus

1. Note that there can be no cases where only the builder country of the vessel changes. By definition, changes in the location of shipyards will always imply that the vessel differs as well.

2. Let $P(V)$ denote the probability of a vessel change. The within- and across-carrier vessel switches are given by $P(V, \neg C)$ and $P(V, C)$, respectively. From Table 1, $P(V) = 0.699$, $P(V, \neg C) = 0.487$, and $P(V, C) = 0.212$. Hence, $P(V, \neg C)/P(V) = 0.487/0.699 \approx 0.70$ and $P(V, C)/P(V) = 0.212/0.699 \approx 0.30$.

Figure 1 — Concentration of shipments of individual products across ship types
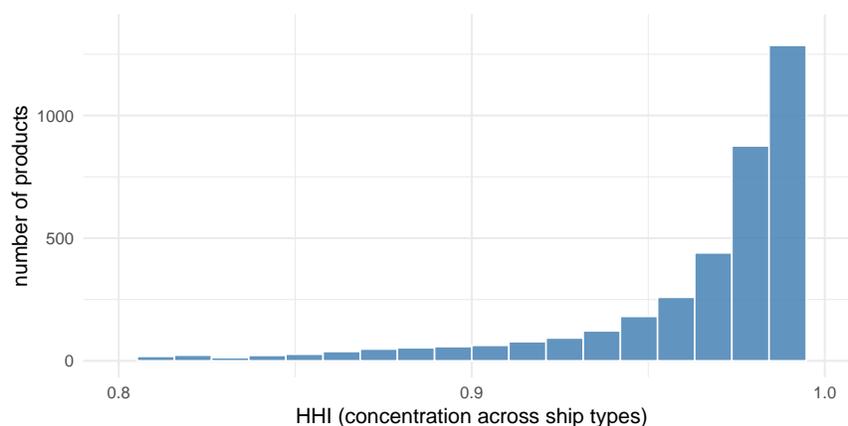
Note: This figure plots the Herfindahl–Hirschman Index (HHI) for each HS-6 product, based on the distribution of its US maritime import shipments across various ship groups during 2020–2024. Ship types include container vessels, dry bulkers, tankers, general cargo, Ro-Ros, and other specialized vessels.

fluctuations in vessel use is driven by carriers reallocating ships differently over time rather than importing firms switching their transport service providers.

Carrier choices are seen to be stickier, indicating the presence of relationship-specific fixed costs. Table 1 further suggests that route choices are relatively persistent, with transactions rarely changing their entry point for goods once sourcing decisions are made. This is seen by the low probability of changing the US port of unlading for a given firm, origin country and product combination (6.6%).

Taken together, the decompositions reveal a clear hierarchy of flexibility within global shipping networks. First, the equipment margin appears to be relatively elastic as vessels frequently differ across firm-level shipments. At an intermediate level, organizational adjustments such as switching carriers or vessel builders occur less frequently. Finally, the spatial dimension is the least flexible as importers rarely switch between US ports.

Stylized fact 2. **Majority of vessel substitutions occur between ships of the same type and size class**

Given the high rate of vessel turnover across shipments, we now turn towards characterizing the nature of substitutability across vessel types. First, we find that maritime transport is highly product-specific. Most products map uniquely to a single ship category.[3] While this tight technological matching is well understood in naval architecture, it has not previously been documented

---

3. Ship categories are broadly defined as container vessels, dry bulkers, tankers, general cargo ships, Ro-Ros and specialised vessels such as nuclear fuel carriers.

Figure 2 — Transitions across size classes conditional on vessel changes
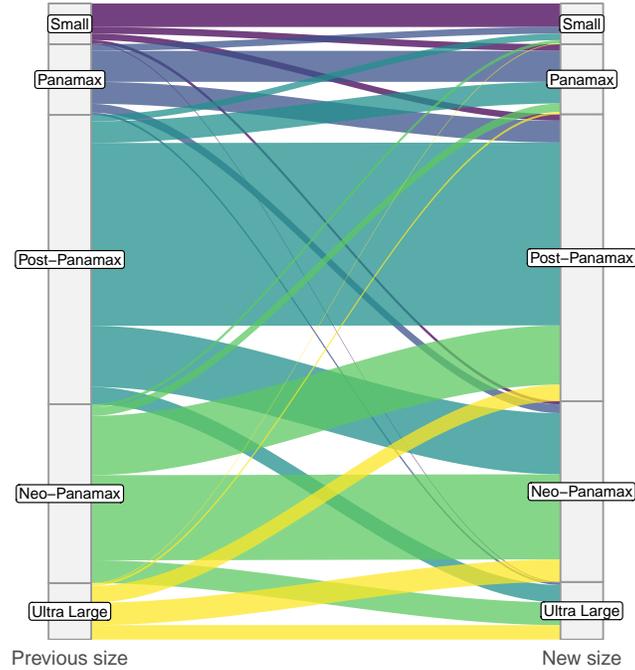
Note: This figure reports transition probabilities across container ship size classes when carriers switch vessels between consecutive shipments of the same US importer–product–origin country triplet over 2020-2024. Size classes are defined based on standard TEU thresholds.

in trade data since (i) AIS databases lack product-level information and (ii) customs data typically omit vessel identifiers.

Figure 1 reveals this correspondence by plotting the Herfindahl index of ship type usage for each product, using bill of lading data on US maritime imports from 2020–2024. The distribution indicates that shipments of products are highly concentrated within specific ship types, implying that substitutions between vessels occur within rather than across ship categories.[4]

Conditional on switching vessels, we next examine whether carriers substitute across size classes for consecutive shipments of a given importing firm-product-origin country combination. The dominant pattern is persistence within the same size class. Figure 2 illustrates this for container vessels, showing the transition probabilities between size classes.[5] In 53% of vessel switches, the

---

4. Cases where a product appears on different ship types arise because some vessels have additional facilities. For example, container or general cargo ships can be equipped with Ro-Ro ramps which allow them to transport vehicles as well.

5. Size classes are defined according to standard TEU thresholds as follows: Small (<3,000 TEU), Panamax (3,000–5,000), Post-Panamax (5,000–10,000), Neo-Panamax (10,000–14,000), Ultra Large (>14,000).

Figure 3 — Builder country shares in newly built vessels
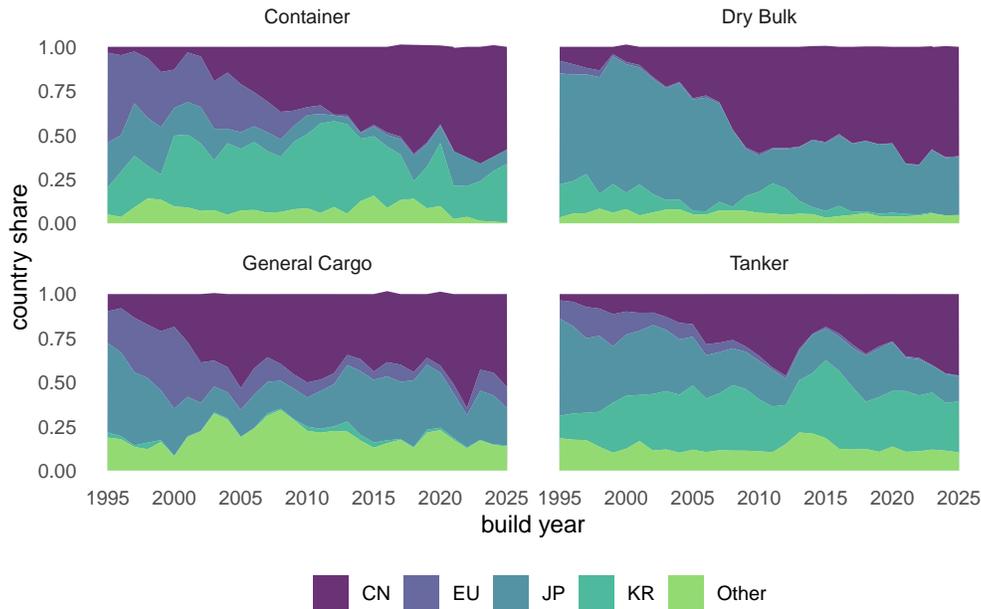
Note: This figure reports builder shares in newly constructed vessels, disaggregated by ship type and build year. Calculations are based on all internationally registered ships built over 1995-2025 which account for more than 75% of the global fleet.

replacement vessel remains in the same size class. A further 35% involve a move to an adjacent class, while only 12% constitute jumps across the size distribution. The pattern extends to other ship categories. Dry bulk carriers remain in the same size class in 65% of vessel replacements. Tankers exhibit greater persistence, with 82% of switches occurring within the same size.[6] These findings indicate that carriers seldom overhaul the vessel size used to move a given origin-product combination. This is consistent with the stickiness we documented in port choices previously. As most origin–product pairs arrive at a given US port, the feasible set of vessel size classes is likely determined by the operational characteristics of those ports.

Taken together, the evidence shows that the maritime transport is segmented by technology. Most products rely on a single ship type and most origin-product pairs are imported via a specific range of vessel sizes. We next investigate how this technological segmentation maps into the shipbuilding industry. To do so, we document the degree of specialization of shipbuilding countries across vessel types and size classes.

Stylized fact 3. **Shipbuilding capacity is concentrated with national specialization across vessel types and size classes**

---

6. See Table A1 in Appendix A.3 for a comparison of transition probabilities across sizes for these vessel types.

Global shipbuilding is a concentrated market and has undergone substantial geographic reallocation over the past three decades (Figure 3). In the mid-1990s, Japan dominated the construction of new dry bulkers and tankers, European shipyards led in container and general cargo vessels, and South Korea occupied a sizable share across fleet segments. China's position was marginal, with lower than 10% share of new builds in any ship category. By 2025, the market shifted dramatically.

China now accounts for more than half of global production in every major vessel class, including dry bulk (62%), containers (58%), general cargo (53%), and tankers (46%). These gains have come with collapsing shares of shipyards in Japan and Europe. South Korea remains competitive, particularly in containers and tankers, but is second to China in terms of the newest ships. The result is a shipbuilding industry that is both more concentrated compared to 1995 (in terms of the Herfindal index of newly built vessels) and reoriented towards China.

When looking at the global stock of vessels, we find that builder countries occupy markedly different positions across various segments of the fleet. South Korea dominates in capital intensive large-scale vessels such as ultra-large container ships, Capesize bulkers, and VLCC-class tankers. European builders are important only in smaller general cargo vessels. China produces across all ship categories and is a key supplier of VLOC dry bulkers and Handysize general cargo vessels, but has not yet displaced South Korea in the highest segments of tankers and container vessels.

Stylized fact 4. **China-built vessels are extensively used across importing firms, industries, source countries, shipping routes and carriers**

Our next stylized fact documents the prevalent use of China-built vessels. Here, we find that Chinese vessels represent one of the most important classes of ships serving the US. In 2024, nearly one in four US importing firms received at least one shipment through a Chinese vessel. Among importers served by more than the mean number of vessels (17 over the year), the average firm was connected to around ten different Chinese vessels. Chinese vessels also accounted for 28% of all ships arriving at US ports in 2024. Only South Korea-built vessels occupy a greater share in US maritime imports, serving 38% of US importing firms and capturing 33% of all vessels arriving in 2024.[7]

Across sectors, US agri-food and dry bulk imports are particularly intensive in their use of Chinese ships. More than 25% of total US maritime imports of products such as sugar, wheat, vegetables, fruit, wool and ferrous metals were through these vessels over 2020-2024. Chinese vessel share is lower but still significant (above 15%) when looking at top US import sectors such as motor vehicles, vehicle parts and other manufactures.

---

7. For shares in US imports across vessels from different builder countries, see Figure A1.

## Figure 4 — Builder shares across different segments of the global fleet
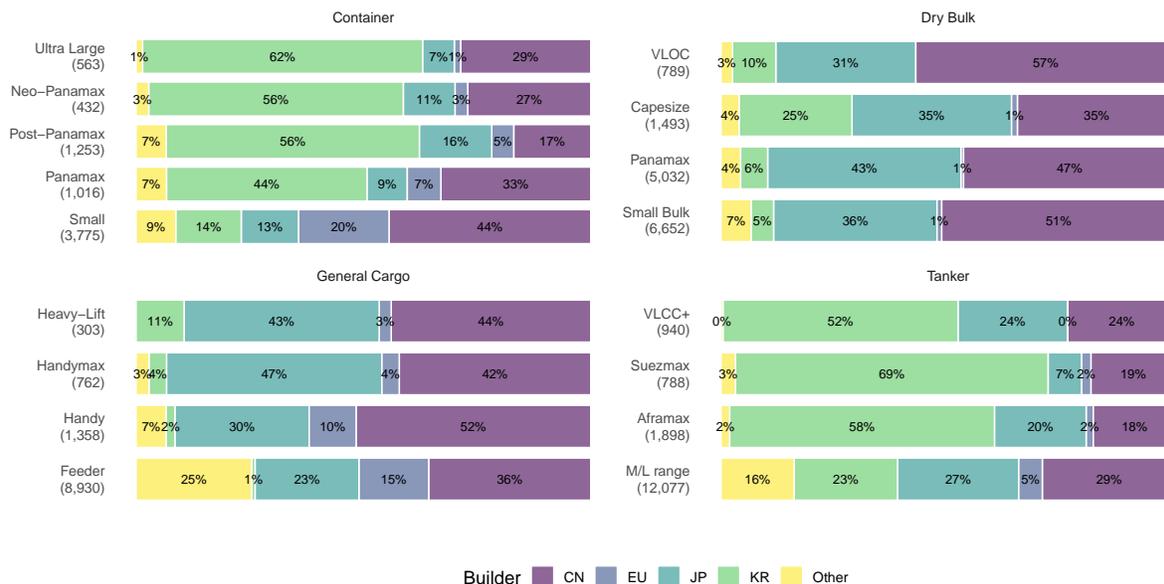
Note: This figure reports builder shares within each ship type–size class segment of the global fleet. Bars represent within-segment shares and size class labels report the total number of vessels in that segment. Calculations are based on all internationally registered ships built over 1995-2025.

Reliance on Chinese-built ships extends well beyond direct trade with China as well. Figure 5 illustrates this exposure across countries, ranking them by the share of maritime exports to the US that is transported via China-built vessels. In the top decile are countries such as Mozambique, Costa Rica, Bolivia and Croatia for whom more than 40% of maritime exports to the US are carried by Chinese-built ships.

Globally, Chinese-built ships account for a substantial fraction of port calls across multiple trade corridors. Figure A2 in the Appendix shows the share of China-built vessels in all vessel port calls across and within regions. Exposure is particularly high in South–South and emerging market corridors. In container shipping, over 48% of all port calls from the Middle East & North Africa (MENA) to Sub-Saharan Africa (SSA) were made by China-built vessels. China's share is similarly high for container ship port calls between Latin America and the Caribbean (LAC) economies (42%). These results reflect broader integration of Chinese shipbuilding into the global merchant fleet, with relatively high exposure on developing region trade flows.

Finally, all major global carriers deploy China-built vessels as part of their operating fleets (Figure A3 in Appendix). China-based COSCO and France-based CMA CGM exhibit the highest reliance, with 56% and 43% of their fleets built in China, respectively. The world's largest carrier MSC that is headquartered in Switzerland, also operates a substantial share of China-built ships with

Figure 5 — Exposure to China-built ships in maritime exports to the US
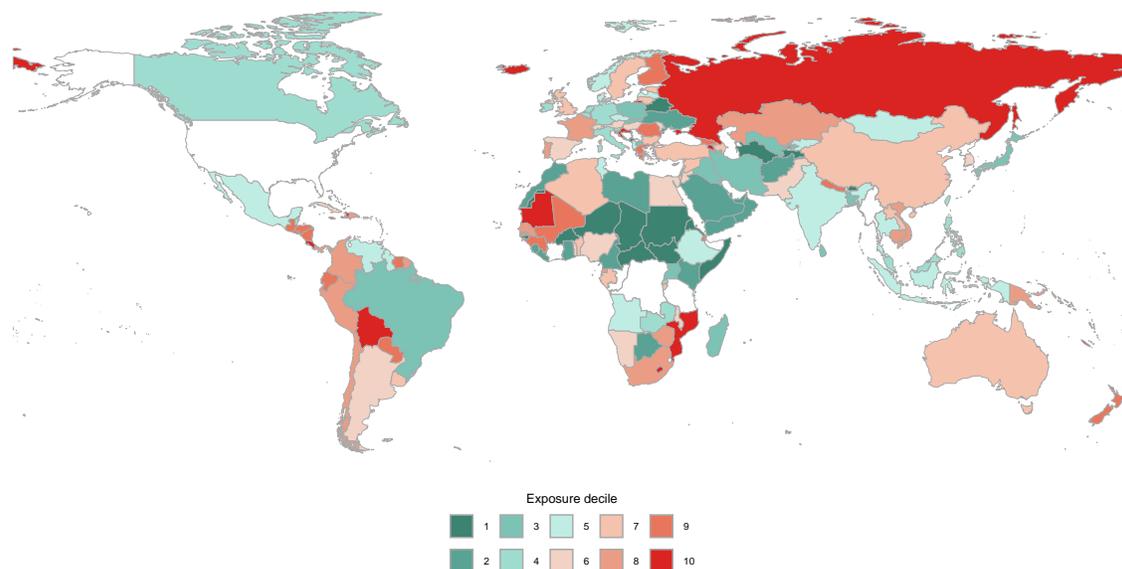
Note: This figure maps countries' exposure to China-built vessels in their maritime exports to the US in 2024. Exposure is measured as the share of US maritime import value transported via ships built in China for each origin country. Given the skewness of exposure levels, countries are assigned to deciles of this distribution.

roughly 27% of its fleet constructed in Chinese shipyards.

Overall, this evidence suggests that imposing port fees on China-built vessels could raise shipping costs not only for trade with China but also for US imports from third countries that rely on Chinese-built ships. Moreover, the reallocation of these vessels across international routes could generate spillover effects by changing transport costs between other country pairs. The effects of such port fees is also likely to hit import sectors and carriers differentially depending on their fleet composition.

4. Estimating the elasticity of substitution across vessels

Having documented these stylized facts, we next turn towards estimating the elasticity of substitution across vessels. To estimate this elasticity, we require an exogenous shock that differen-

tially affects transportation costs across vessel types. A promising source of variation arises from large-scale port infrastructure investments that have changed transportation costs differentially by vessel size. In particular, we exploit the 2016 expansion of the Panama Canal which introduced new locks capable of accommodating substantially larger ships. This reform potentially reduced the relative cost of deploying larger vessels on routes exposed to the Panama Canal. We therefore leverage this quasi-natural experiment along with cross-sectional variation in vessel size, builder country, and emissions intensity to identify the implied substitution patterns across vessel types.

4.1. The Panama Canal expansion of 2016

Situated at the narrowest point between the Atlantic and Pacific Oceans, the Panama Canal bypasses the need for ships to undertake lengthy passages around Cape Horn. By the early 2000s however, sustained growth in world trade and advances in shipbuilding had rendered the Canal's original locks increasingly restrictive. Traffic volumes had risen beyond the infrastructure's initial design capacity, and a growing share of new container and bulk vessels exceeded the Canal's size limits.

Facing these constraints, the Panama Canal Authority initiated a major expansion project in 2007 to construct a new, deeper, and wider lane capable of accommodating the next generation of large ships. Approved through a national referendum and budgeted at approximately USD 5.25 billion, the expansion added a third traffic lane equipped with two new lock complexes (Agua Clara on the Atlantic and Cocolí on the Pacific side) which opened simultaneously to commercial traffic in June 2016.

This expansion offers a particularly clean setting to study substitution across vessels. First, it constitutes a significant upgrading of infrastructure in a major artery of the global shipping network. Existing research confirms that the 2016 expansion induced substantial trade effects. Shipping costs declined by 2-3% along treated routes, defined as port pairs whose fastest connection passes via the Panama Canal using observed AIS vessel movements (Heiland et al. 2025). This differential exposure of routes provides an important source of variation which enables us to identify substitution patterns across vessels deployed on treated relative to control routes.

Second, the Canal expansion set clear dimensional thresholds on vessel length, beam, and draft which could pass through the new locks. These thresholds create a well-defined treatment boundary across vessels. Ships exceeding the limits of the old locks but within the new size constraints suddenly gained access to the Canal, while those above the new cutoff remained excluded. This discontinuous expansion of eligibility generates a transparent distinction between treated and untreated vessel classes, which we exploit to identify differential responses to the expansion.

Finally, the Canal's expansion is particularly informative in the US context. Trade routes linking US East Coast ports with major Asian and Latin American partners were exposed and likely experienced first-order reductions in transportation costs. In the case of the US, we also observe detailed bill of lading records that allow us to track shipments and vessel deployment at high frequency, enabling us to study the dynamics of how the composition of vessel fleets changed on these routes due to the Canal's upgraded capacity.

4.2. Empirical specification

To estimate the elasticity of substitution across vessel types, we implement a Wald difference-in-differences (Wald-DiD) estimator. The Wald-DiD identifies an elasticity as the ratio of two reduced-form estimates, namely, the effect of the policy shock on relative quantities and its effect on relative prices or costs. In our setting, the expansion of the Panama Canal serves as a transportation cost shock with differential effects across vessels depending on their size and across routes depending on their reliance on the Canal.

We use this variation to construct two components of the Wald-DiD: (i) the relative change in import volumes carried by large versus smaller vessels, and (ii) the relative change in transport costs along treated versus control routes. The ratio of these two effects yields the implied elasticity of substitution across vessel types. This ratio can also be interpreted as a two-stage least squares (2SLS) estimator, where the first stage captures the impact of the Panama Canal expansion on transportation costs, and the second stage measures the impact on relative trade flows by vessel size class. Our baseline event-study specification is given by equation (1).

$$\text{Y}_{opt} = \sum_{k=-6}^{6} \beta_k \cdot \mathbf{1}\{\text{event time}_{opt} = k\} \times \text{Panama route}_{op} + \theta_{op} + \theta_{ot} + \theta_{pt} + \varepsilon_{opt} \qquad (1)$$

In this specification, o indexes the origin country, p the US port of unlading, and t the year-month. The dependent variable Y_{opt} captures different dimensions of adjustment in maritime trade flows. We consider three outcomes. First, the log percentile of vessel capacity (TEU) which reflects the size distribution of the container fleet deployed within origin country - US port pairs. Second, the share of US maritime imports carried by each vessel size class for a given origin–port pair, which captures changes in their relative market shares. Third, the average maritime freight charge, which serves as a proxy for transportation costs. Together, these outcomes allow us to disentangle adjustments in fleet composition, trade allocation across vessel size classes and changes in shipping costs.

Treated routes are defined by Panama route$_{op}$ which is an indicator equal to one for combinations of US East Coast ports and specific origin countries that are likely exposed to the Panama Canal.[8] This indicator is interacted with a series of event-time dummies that track months relative to the Canal's expansion. Here, we take the month corresponding to the opening of new locks as the reference period (June 2016).

Equation (1) further includes a comprehensive set of fixed effects to reduce risk of omitted variable bias. Origin–port fixed effects (θ_{op}) absorb all time-invariant determinants of trade between a given country and US port, such as geographical location, historical trade links and logistical complementarities. Origin–time fixed effects (θ_{ot}) control for macroeconomic conditions in exporting countries that can impact trade flows independent of the Canal expansion. Finally, port–time fixed effects (θ_{pt}) capture demand and infrastructure-related shocks at the US port level, including local congestion or dock labor disruptions. With these fixed effects, identification comes from time variation within origin country and US port pairs. Standard errors are clustered by all dimensions, o, p and t.

4.3. Event study results

Following the specification described above, the results reveal a clear and immediate shift in the composition of vessels operating on treated routes. Figure 6 plots the estimated β_k coefficients from equation (1) for the six months before and after the Canal expansion. The figure shows changes in the 25th, 50th, and 75th percentiles of vessel sizes (measured by total TEU) for container ships transporting goods from each origin country to each US port over time.

The results indicate a pronounced rightward shift in the vessel size distribution. Within six months of the Canal's expansion, vessel sizes on treated routes rose sharply with increases of roughly 13.4% at the 25th percentile, 14.2% at the median, and 9.5% at the 75th percentile. Overall, the pattern points to a broad-based upgrading of the fleet rather than the entry of a few exceptionally large ships. The speed of adjustment further suggests that the new locks effectively removed a binding capacity constraint on these routes.

Next, we examine how the allocation of US maritime import volumes across vessel size categories evolved. Before the expansion, the largest ships able to transit the Canal were Panamax vessels (3,000-5,000 TEU). The new locks enabled Post-Panamax vessels (5,000–10,000 TEU) to operate on these routes. We therefore track how the shares of total US seaborne imports (by volume) transported by these vessel categories changed before and after the expansion (Figure 7).

8. Origin countries considered here are East Asian economies (China, Japan, South Korea, Taiwan), South East Asian economies (Vietnam, Thailand, Malaysia, Indonesia, Singapore, Philippines and Cambodia), Oceania (Australia and New Zealand) and the Western coast of South America (Chile, Peru, Ecuador and Colombia).

Figure 6 — Changes in vessel size distribution following Panama Canal expansion
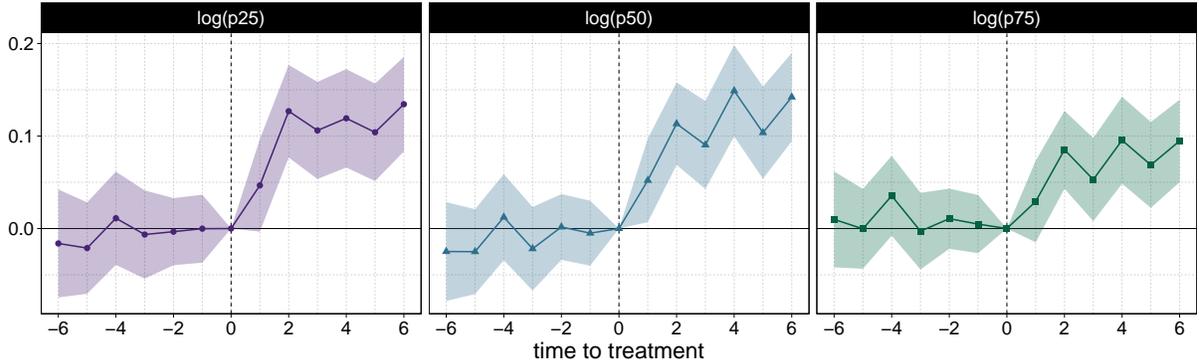
Note: This figure reports the estimated β_k coefficients from equation (1) for the six months before and after the Panama Canal expansion in 2016. The coefficients trace changes in the log 25th, 50th and 75th percentiles of the vessel size distribution, for all container ships deployed between different origin country and US port combinations.

The results show that the share of Panamax vessels begins to decline sharply from June 2016, falling by roughly 18 pp within six months. This indicates that Panamax ships, which had previously been the largest that could transit the Canal, were gradually displaced from treated routes once the new locks opened. On the right panel, the share of Post-Panamax vessels rises after the expansion. By December 2016, its share had risen by more than 13 pp relative to control routes. These two panels together suggest a substantial reallocation of traffic from Panamax to larger Post-Panamax ships, consistent with the removal of the Canal's capacity constraint.[9] These results are also robust to using import values instead of import volumes (see Figure A6).

This shift toward larger vessels provides the first component of the Wald-DiD estimator. These findings are also confirmed by industry reports which indicate that the Canal's opening induced a collapse in the market for Panamax vessels that were previously a workhorse of the container ship industry.

To ensure that these patterns are not driven by unrelated shifts in global shipping or measurement noise, we next conduct a placebo test. The Panama Canal expansion defined thresholds for vessel eligibility which create clear treatment boundaries by ship size. We now exploit this feature to examine vessel types that should have been unaffected by the expansion. Specifically, we consider smaller vessels (below 3,000 TEU) which were already able to transit the Canal prior to the expansion and ultra-large vessels (above 14,000 TEU) which remained too large to transit even after the new locks opened. Figure 8 reports the corresponding estimates from equation

---

9. In addition to the growth of Post-Panamax shipping, we also observe modest and more gradual reallocation of US import volumes towards Neo-Panamax vessels which range from 10,000 - 14,000 TEU. See Figure A4 in Appendix A.3 for further details.
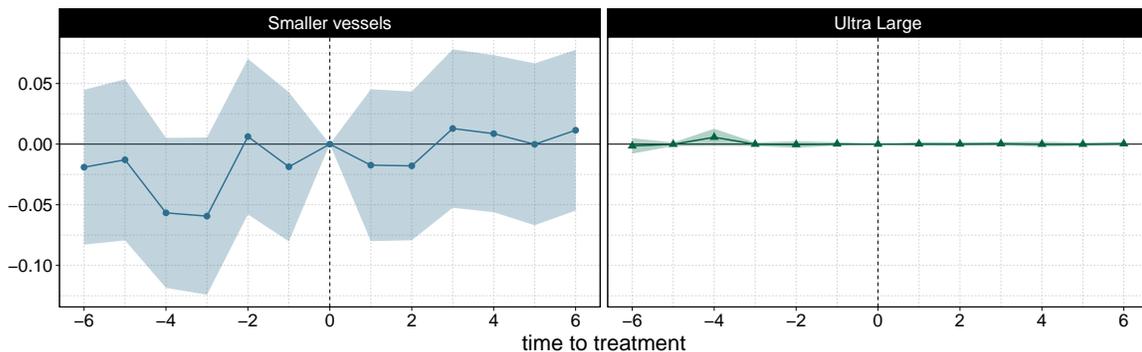
Figure 7 — Shares of Panamax and Post-Panamax in US import volumes
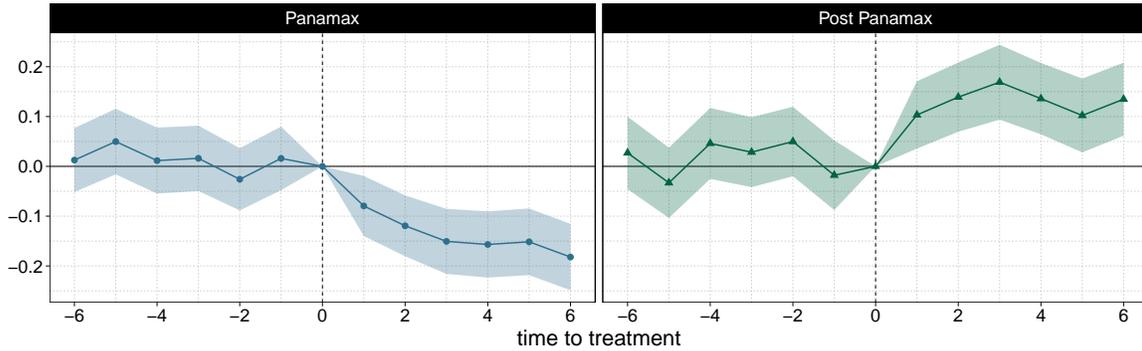
Note: This figure reports estimated β_k coefficients from equation (1) for the six months before and after the Panama Canal expansion in 2016. The dependent variables correspond to the share of total US maritime import volume that is transported via Panamax (left) and Post-Panamax container ships (right).

(1), where the dependent variable measures the respective shares of these two vessel categories in US import volumes. In line with our priors, there is no significant change in trade volumes transported by small and ultra-large vessels. Thus the major reshuffling of the fleet occurred in the middle of the size distribution with the switch from Panamax to Post-Panamax vessels.

Figure 8 — Shares of small and ultra-large vessels in US import volumes

Note: This figure reports estimated β_k coefficients from equation (1) for the six months before and after the Panama Canal expansion in 2016. The dependent variables correspond to the share of total US maritime import volume that is transported via small (left) and ultra-large container ships (right).

These results remain robust even when vessels are classified by width rather than their TEU capacity. Figure A5 in the Appendix presents the corresponding estimates. We observe a sharp decline in volumes transported through the Panama Canal on vessels that were always eligible to

18

Figure 9 — Impact of Panama Canal expansion on average freight rate
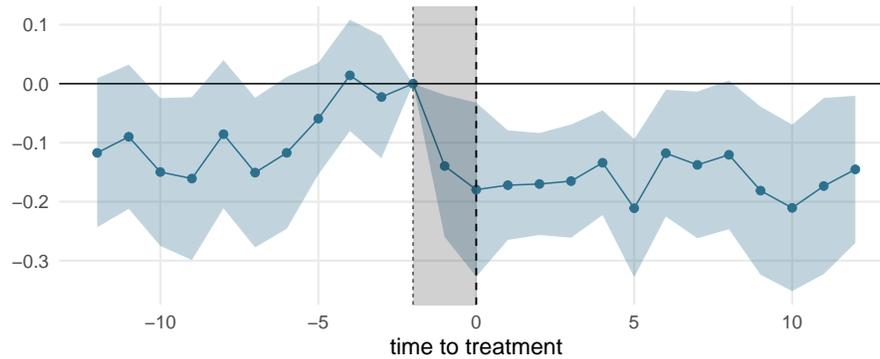
Note: This figure reports estimated β_k coefficients from a modified equation (1), for the 12 months before and after the Panama Canal expansion in 2016. The dependent variable correspond to the log average freight rate per dollar, computed as the ratio of total vessel charges to maritime import values for a given US port district, product and origin country combination.

transit (width below 32.3 meters), mirrored by an increase in the share of import volume carried on newly eligible vessels (width ranging from 32.3 to 49 meters). Together, these findings confirm that the expansion induced a swift reallocation toward ships that gained access to the new locks.

For the second component of the Wald-DiD estimator, we quantify how the Canal expansion affected maritime freight costs. We draw on US Census data, which report monthly values of seaborne imports and cumulative vessel charges by product, origin country, and US port district of unlading. Dividing total vessel charges by the corresponding import values yields an average freight rate per dollar, as a measure of shipping costs that is comparable across routes and over time. We compute these freight rates for the year before and after the Canal expansion and compare treated and control routes defined as in the preceding analysis. Figure 9 displays the estimated effects of the Canal expansion on these freight rates.[10]

The event study estimates show a reduction in maritime freight charges on origin country–US port pairs exposed to the Panama Canal expansion. Interestingly, freight charges declined by approximately 13% one month before the new locks opened. This may be due to carriers adjusting pricing in anticipation of the increase in capacity. An alternative explanation is that exporters may have delayed shipments till the opening of the new locks, reducing short-term demand for freight services and incentivizing carriers to lower prices earlier. When the new locks opened

---

10. Note that the US Census reports vessel charges only at the port-district level rather than for individual ports. As a result, these estimates are not port-specific. However, since ports within a district are typically administered by a single regional authority (for example, the Port Authority of New York and New Jersey), we assume that freight rate variation across ports within the same district is limited.

Table 2 — Combined effect of Panama Canal 2016 expansion

Dependent Variables:	log(relative import value)	log(freight rate)	log(relative import value)
IV stages		First	Second
Model:	(1)	(2)	(3)
exposed route x post expansion	1.069***	-0.180***	
	(0.226)	(0.036)	
log(freight rate)			-5.940***
			(1.831)
origin country x US port district	Yes	Yes	Yes
origin country x year	Yes	Yes	Yes
US port district x year	Yes	Yes	Yes
Observations	23,608	23,608	23,608
R^2	0.608	0.601	–
Cragg-Donald			20.552

Note: This table reports DiD and IV estimates on the impact of the Panama Canal expansion of 2016, at the origin country, US port district and month level over 2015 to 2017. In column (1), the dependent variable is log relative import value, defined as the ratio of US maritime imports carried by Post-Panamax and Neo-Panamax vessels to those carried by Panamax and smaller vessels. The dependent variable in column (2) corresponds to log average freight rate, computed as total vessel charges divided by import values. Column (3) reports the implied elasticity of substitution, where log average freight rates are instrumented with countries' pre-expansion share of maritime exports to the US that transits through the Panama Canal (exposed route). Post expansion is a dummy variable that takes the value of one from June 2016 onwards.

in June 2016, freight charges on treated routes were 18% lower relative to control routes and the pre-expansion period. These costs also remained persistently lower over the subsequent year. Together, these results indicate that the Canal expansion led to a lasting reduction in transport costs by enabling the use of larger vessels. This decline in freight charges also provides the second component of the Wald-DiD estimator.

In the final step toward estimating the elasticity of substitution, we merge bill of lading records with US Census data at the origin country, US port district and monthly level from 2015 to 2017. This allows us to link adjustments in fleet composition to changes in transportation costs for goods sourced from countries that are differentially exposed to the Panama Canal when exporting to the US. The elasticity of substitution is then obtained via a Wald-DiD setup and computed as the ratio of the change in relative import values between larger and smaller vessels to the change in relative freight rates on these routes.[11]

11. The relative import value is defined as the ratio of imports carried by Post-Panamax and Neo-Panamax vessels to those carried by Panamax and smaller vessels, aggregated by origin country, US port district and month.

Formally, this is equivalent to an IV regression in which changes to freight rates are instrumented with countries' reliance on the Panama Canal for shipping to the US. Results from these estimations are shown in Table 2.[12] The estimate in column (3) provides the elasticity of substitution. It shows that a 1% reduction in relative freight charges from the Canal expansion induced an increase of approximately 5.9% in the share of large vessels in exporting to the US.

To our knowledge, this is the first direct estimate of substitution elasticity across vessel size classes in global shipping. Existing work has focused on elasticities across routes or modes of transport, but not within the maritime fleet itself. The relatively high elasticity further suggests that the composition of the fleet adjusts rapidly to cost changes, such that investments in port infrastructure can have significant and immediate effects on transportation choices.

5. A quantitative trade model with vessel choice

Our next step is to incorporate this empirical evidence into a quantitative trade model. To do so, we extend the framework of Caliendo and Parro (2015) who develop a multi-sector version of Eaton and Kortum (2002) with IO linkages. Our contribution here is to further augment this framework with a detailed description of the transportation sector, featuring multiple transport modes (air, land, sea) alongside different categories of vessels. This enables us to conduct counterfactual analyses of policies targeting specific segments of the maritime transportation sector in a quantitative model with multiple countries, sectors, global input output linkages, and a market for vessel transportation services. In the first part of this section, we present a redux of the Caliendo-Parro model, focusing on the relevant mechanisms.[13] In the second part, we describe in detail our novel model of the transportation sector.

5.1. Production and trade

The model describes the trade and production of N countries, indexed o and d, and J sectors, indexed j and k. Production uses labor as the sole factor, which is mobile across sectors but not across countries. All markets are perfectly competitive. Sectors are either wholly tradable or non-tradable.

---

12. Instead of the binary treatment indicator Panama route$_{op}$ in equation (1), Table 2 uses a continuous measure of exposure that captures countries' intensity of Panama Canal use when exporting to the US. To do so, we use bill of lading data from 2014 for the same set of countries that potentially use the Panama Canal when exporting to the US (China, Japan, South Korea, Taiwan, Vietnam, Thailand, Malaysia, Indonesia, Singapore, Philippines, Cambodia, Australia, New Zealand, Chile, Peru, Ecuador and Colombia). For these countries, we calculate the share of exports to US East Coast in their total maritime exports to the US. This continuous measure provides additional cross-country variation in Panama Canal use that strengthens identification in the IV estimations.

13. Section A.1 in the Appendix lays out the model in full detail.

Every country produces final and intermediate goods using domestic and imported varieties of J differentiated goods from all other countries.[14] Intermediate goods production also uses labor. Let X_j^d country d's total expenditure on varieties of good j. Then, country d's imports of sector-j varieties from country o are given by

$$X_{od}^j = \pi_{od}^j X_d^j, \qquad \text{where} \qquad \pi_{od}^j = \frac{\lambda_o^j (c_o^j \phi_{od}^j)^{-\theta^j}}{\sum_{h=1}^N \lambda_h^j (c_h^j \phi_{hd}^j)^{-\theta^j}} \tag{2}$$

equals the share of country j's total expenditure devoted to varieties from sector j in country o. Bilateral trade flows $\pi_{od}^j X_d^j$ follow a sectoral gravity equation. Exports from o to d in sector j depend on the size of the destination market captured by X_d^j, and the relative competitiveness of o as a source country, captured by π_{od}^j. In this Ricardian world with perfect competition, competitiveness is entirely determined by cost. The cost of serving market d faced by a representative firm from country o's sector j depends on trade frictions ϕ_{od}^j, input prices c_o^j, and an inverse measure of average productivity λ_o^j. The iceberg-type trade frictions ϕ_{od}^j combine political barriers to trade and transportation cost. The latter will be derived below from a model of the multimodal transportation sector.

Input prices contain wages w_o and the prices of intermediate inputs combined in the following cost function

$$c_d^j = \frac{\Upsilon_d^j w_d^{\beta_d^j}}{z_d^j(\omega^j)} \left[\prod_{k=1}^J (P_d^k)^{\gamma_d^{k,j}} \right]^{1-\beta_d^j},$$

where P_d^k is the price of a composite intermediate good from sector k and Υ_d^j is a constant. γ_{jk}^d is the share of intermediate goods expenditure sector-j producers spend on the good from sector k, and β_d^j is cost share of labor.

The price of the composite sector-j good in country d is then

$$P_d^j = A^j \left[\sum_{o=1}^N \lambda_o^j (c_o^j \phi_{od}^j)^{-\theta^j} \right]^{-1/\theta^j}, \tag{3}$$

where $A^j = \Gamma(1 + (1-\sigma^j)/\theta^j)^{1/(1-\sigma^j)}$.

---

14. Our sectors span both goods and services. For simplicity, we refer to all of them as goods as long as the distinction is immaterial.

Total expenditures on goods from sector j in country d are given by

$$X_d^j = \sum_{k=1}^{J}(1 - \beta_d^k)\gamma_d^{j,k}Y_d^j + \alpha_d^j I_d. \tag{4}$$

The first term on the right-hand side is the expenditure on intermediate inputs of type j, a share $(1 - \beta_d^k)\gamma_d^{j,k}$ of each sector k's production value Y_k. The second term denotes final expenditure on goods from sector j, given by a constant share α_d^j of country d's income I_d. Sectoral goods market clearing $Y_d^j = \sum_{o=1}^{N} X_o^k \frac{\pi_{do}^k}{\phi_{do}^k}$ pins down expenditure levels as functions of the above trade patterns and wages.

The final step towards general equilibrium is to determine wages. To that end, we invoke an income-equals-expenditure condition. Aggregate income in country d is

$$I_d = w_d L_d + R_d + D_d,$$

where $w_d L_d$ is labor income and R_d is net revenue from trade and transport policies plus redistributed profits from the shipping sector, which we describe in more detail below. D_d denotes the aggregate trade deficit, which, in our static model, reflects an exogenously given transfer to or from the rest of the world.

The trade deficit then obtains as aggregate imports (including domestic purchases) minus aggregate exports (including domestic sales)

$$D_d = \sum_{j=1}^{J}\sum_{o=1}^{N} X_{od}^j - \sum_{j=1}^{J}\sum_{o=1}^{N} X_{do}^j, \tag{5}$$

5.2. Multiple transport modes and vessel choice

We now turn to our model extension describing the international transportation sector, where the transportation cost of tradable shipments are (partly) determined endogenously. Similar to Tolva (2024), we assume that for every shipment of tradable goods between a pair of countries (o, d), a transport mode is chosen

$$m \in \mathcal{M} = \{\text{air}, \text{sea}, \text{land}\},$$

23

Conditional on selecting *sea*—one of multiple vessel types is selected. We structure the vessel-type choice as follows. First, we assume that there is one dominant vessel type

$$\kappa^j \in \mathcal{K} = \{Container, Tanker, Bulk\}$$

for each tradable sector. Second, within the sector-specific vessel type, a size class (large or small) $s \in \mathcal{V} = \{L, S\}$ is chosen, and, conditional on size, a specific builder country $z \in \mathcal{Z}$. We assume that these choices follow a nested CES (multinomial–logit) function with three nests.

(i) Builder country nest (inside size), elasticity $\zeta > 1$.

Let $d_{od}^{\text{sea},\kappa^j,s,z,j}$ denote the delivered cost of shipping sector j goods from o to d on a vessel of size s built in country z. The cost of transportation by vessel type are endogenously determined in a market for ship services, which we will describe further below.

The composite cost of vessels of size s in sector j's optimal ship class κ^j is then

$$d_{od}^{\text{sea}\kappa^j,s,j} = \left[\sum_{z \in \mathcal{Z}} (d_{od}^{\text{sea},\kappa^j,s,z,j})^{1-\zeta} \right]^{1/(1-\zeta)},$$

with conditional builder-country shares

$$\tilde{\pi}_{od}^{\kappa^j,z,s,j} = \frac{\left(d_{od}^{\text{sea},\kappa^j,s,z} \right)^{1-\zeta}}{\sum_{u \in \mathcal{Z}} (d_{od}^{\text{sea},\kappa^j,s,u})^{1-\zeta}}, \qquad z \in \mathcal{Z}.$$

(ii) Vessel size nest (inside sea), elasticity $\rho > 1$. The composite cost of the sea leg for sector j is

$$d_{od}^{\text{sea},\kappa^j,j} = \left[\sum_{s \in \mathcal{S}} (d_{od}^{\text{sea},\kappa^j,s,j})^{1-\rho} \right]^{1/(1-\rho)},$$

with conditional size-class shares

$$\tilde{\pi}_{od}^{\kappa^j,s,j} = \frac{\left(d_{od}^{\text{sea},\kappa^j,s,j} \right)^{1-\rho}}{\sum_{r \in \mathcal{S}} (d_{od}^{\text{sea},\kappa^j,r,j})^{1-\rho}}, \qquad s \in \mathcal{S}.$$

(iii) Mode nest, elasticity $\eta > 1$. The composite transport friction for sector j between (o, d) is

$$d_{od}^j = \left[(d_{od}^{\text{air},j})^{1-\eta} + (d_{od}^{\text{land},j})^{1-\eta} + (d_{od}^{\text{sea},j})^{1-\eta} \right]^{1/(1-\eta)},$$

with mode shares

$$\pi_{od}^{m,j} = \frac{(d_{od}^{m,j})^{1-\eta}}{\sum_{\tilde{m} \in \mathcal{M}} (d_{od}^{\tilde{m},j})^{1-\eta}}, \qquad m \in \mathcal{M},$$

where $d_{od}^{air,j}$ and $d_{od}^{land,j}$ denote the delivered cost for modes air and land, respectively. We assume that $d_{od}^{air,j}$ and $d_{od}^{land,j}$ are exogenous in our main analysis, but provide a robustness analysis where the price of air and land transportation rises with the amount of goods transported.[15]

Interpretation. The term $\pi_{od}^{m,j}$ is the *mode share*: the fraction of sector-j imports that travels from o to d using mode $m \in \mathcal{M}$. Within the sea leg, $\tilde{\pi}_{od}^{\kappa^j,s,j}$ is the *conditional size-class share*, i.e. the probability that a sea shipment is carried by vessel of size $s \in \mathcal{S}$ in sector j's vessel class κ^j. Conditional on size and class, $\tilde{\pi}_{od}^{\kappa^j,z,s,j}$ is the probability that the shipment is carried by a vessel built in country z. By construction $\sum_m \pi_{od}^{m,j} = 1$ and, if $\pi_{od}^{\text{sea},j} > 0$, then $\sum_s \tilde{\pi}_{od}^{\kappa^j,s,j} = 1$ and $\sum_z \tilde{\pi}_{od}^{\kappa^j,s,z,j} = 1$; otherwise the vessel shares are not defined.

For tradable sectors, the composite transportation cost augment other (exogenously given) barriers to trade $\bar{\phi}_{od}^j$, such that

$$
\phi_{od}^j = \begin{cases}
\bar{\phi}_{od}^j d_{od}^j & \text{for } j \in \mathcal{J}^{\text{G}} \text{ (tradable merchandise)} \\
\bar{\phi}_{od}^j & \text{for } j \in \mathcal{J}^{\text{S}} \text{ (tradable services)} \\
\bar{\phi}_{od}^j \text{ if } o = d \text{ and } \infty \text{ if } o \neq d & \text{for } j \in \mathcal{J}^{\text{N}} \text{ (non-tradable sectors)}
\end{cases}
$$

5.3. Endogeneous determination of vessel service cost

A unit of ship service of type (κ, s, z) is traded on a world market at price $\Xi^{\kappa,s,z}$ per unit value transported. The delivered cost of shipping a good in sector j from o to d using vessel type of this type is then

$$
d_{od}^{\text{sea},\kappa,s,z,j} = \delta_{od}^{\text{sea},\kappa^j,s,z,j} \Xi^{\kappa^j,s,z},
$$

where $\delta_{od}^{\text{sea},v,j}$ is an exogenously given **route–sector–type specific cost shifter**, such as a fee for the usage of vessels from a particular builder country on a particular route.

Let X_{od}^j denote expenditure on sector-j goods imported by d from o. Sea-borne expenditure on vessels of type (κ^j, s, z) for sector j on route (o, d) is then

$$
X_{od}^{\text{sea},\kappa^j,s,z,j} \equiv \tilde{\pi}_{od}^{\text{sea},\kappa^j,s,z,j} \tilde{\pi}_{od}^{\text{sea},s,j} \tilde{\pi}_{od}^{\text{sea},j} X_{od}^j,
$$

The corresponding demand for shipping capacity equals

$$
DWT_{od}^{\text{sea},\kappa^j,s,z,j} = \frac{(d^{sea,\kappa^j,s,z,j})^{-\zeta}}{\sum_{u \in \mathcal{Z}} (d_{od}^{\text{sea},\kappa^j,s,u})^{-\zeta}} \frac{(d^{sea,\kappa^j,s,j})^{-\rho}}{\sum_{r \in \mathcal{S}} (d_{od}^{\text{sea},\kappa^j,r,j})^{-\rho}} \frac{(d^{sea,j})^{-\eta}}{\sum_{\tilde{m} \in \mathcal{M}} (d_{od}^{\tilde{m},j})^{-\eta}} \frac{X_{od}^j}{P_d^j}
$$

---

15. Specifically, we follow Allen and Arkolakis (2022), Tolva (2024), and Fuchs and Wong (2024) and model the delivered-cost of air and land transportation as the product of an exogenous component and a congestion surcharge $d_{od}^{m,j} = \delta_{od}^{m,j} \Xi_{od}^m$, where $\Xi_{od}^m = (X_{od}^m / \bar{X}_{od}^m)^\varphi$ and $0 \leq \varphi < 1$.

Total demand for shipping capacity (deadweight tonnes, DWT) of type (κ, s, z) is then the sum of the demand from all sectors that use vessels of class κ:

$$DWT_{od}^{\kappa,s,z} = \sum_{j \in \mathcal{J}_{\kappa^j=\kappa}^{G}} DWT_{od}^{\text{sea},\kappa^j,s,z,j}$$

For each vessel type, we assume a fixed world supply of ship capacity $DWT^{\kappa,s,z,\text{supply}}$. Ship-service prices Ξ^v adjust to clear the world market:

$$\sum_{o,d} DWT_{od}^{\kappa,s,z} = DWT^{\kappa,s,z,\text{supply}}, \qquad \kappa \in \mathcal{K}, s \in \mathcal{S}, z \in \mathcal{Z}.$$

These market-clearing conditions determine the ship-service prices $(\Xi^{\kappa,s,z})$ given bilateral trade flows, transportation choices, and the baseline route-specific cost shifters $\delta_{od}^{\text{sea},\kappa^j,s,z,j}$. Changes in the demand for ship services on some routes therefore feed back into the prices $\Xi^{\kappa,s,z}$ and propagate to third countries via the global market for shipping services, in addition to the standard general-equilibrium linkages through goods markets.

5.4. Ship service income and policy revenue

We now describe how income from ship services and policy-induced revenue introduced through R_d in the income equation above are determined and distributed across countries.

Ship service income. Ship service providers earn revenue from transporting goods. In the baseline equilibrium, we assume zero profits: ship service prices equal marginal costs, so providers earn exactly what is required to cover costs. The baseline ship service value for vessel type (κ, s, z) is

$$(\Xi^{\kappa,s,z} - 1)\frac{Q^{\kappa,s,z}}{\Xi^{\kappa,s,z}} \qquad \text{where} \qquad Q^{\kappa,s,z} = \sum_{o,d} \sum_{j \in \mathcal{J}_{\kappa^j=\kappa}^{G}} \pi_{od}^{\text{sea},\kappa^j,s,z,j} \cdot X_{od}^j$$

$Q^{\kappa,s,z}$ is the total value of goods transported on vessel type (κ, s, z) and $\pi_{od}^{\text{sea},\kappa^j,s,z,j} = \tilde{\pi}_{od}^{\kappa^j,s,z,j} \cdot \tilde{\pi}_{od}^{\kappa^j,s,j} \cdot \pi_{od}^{\text{sea},j} \cdot \pi_{od}^{j}$ is the share of bilateral sectoral expenditure transported on vessels of type (κ, s, z).

In the counterfactual, when ship service prices change, providers earn revenue that deviates from baseline cost recovery. The profits (or losses) generated in market for ship services are

$$\Delta I^{\text{ship}} = \sum_{\kappa,s,z} \left((\Xi^{\kappa,s,z\prime} - 1) \frac{Q^{\kappa,s,z,\prime}}{\Xi^{\kappa,s,z,\prime}} - (\Xi^{\kappa,s,z} - 1) \frac{Q^{\kappa,s,z}}{\Xi^{\kappa,s,z}} \right)$$

where $Q^{\prime\kappa,s,z}$ is the counterfactual ship service value. When ship prices fall ($\hat{\Xi}^{\kappa,s,z} < 1$), this term is negative, representing a loss to ship service providers.

We assume ship service providers are owned globally in proportion to GDP. The ship service income change accruing to country d is therefore

$$\Delta I_d^{\text{ship}} = \frac{w_d L_d}{\sum_c w_c L_c} \cdot \Delta I^{\text{ship}}.$$

Policy revenue from transport cost wedges. Policy interventions may introduce cost wedges $\tau_{od}^{\kappa,s,z,j} \geq 1$ on specific transport routes. These wedges could represent port fees, carbon taxes, or other levies on maritime transport. The policy-induced cost shifter is

$$\delta_{od}^{\text{sea},\kappa^j,s,z,j} = \tau_{od}^{\kappa,s,z,j} \cdot \bar{\delta}_{od}^{\text{sea},\kappa^j,s,z,j},$$

where $\bar{\delta}_{od}^{\text{sea},\kappa^j,s,z,j}$ represents baseline (non-policy) route characteristics.

The total revenue generated by these wedges is

$$R^{\text{policy}} = \sum_{o,d,\kappa,s,z,j} \frac{\tau_{od}^{\kappa,s,z,j} - 1}{\tau_{od}^{\kappa,s,z,j}} \cdot X_{od}^{\text{sea},\kappa,s,z,j\prime},$$

where $X_{od}^{\text{sea},\kappa,s,z,j\prime}$ is the counterfactual sea-borne expenditure.

The distribution of this revenue depends on the nature of the policy.

(i) Destination-based and/or origin based collection.

Policies targeting both arriving and departing vessels generate revenue on both the import and export side, such that $R_d^{policy} = R_d^{policy,imp} + R_d^{policy,exp}$ with

$$R_d^{\text{policy,imp}} = \sum_{o,\kappa,s,z,j} \frac{\tau_{od}^{\kappa,s,z,j} - 1}{\tau_{od}^{\kappa,s,z,j}} \cdot X_{od}^{\text{sea},\kappa,s,z,j\prime} \qquad \text{and} \qquad R_d^{\text{policy,exp}} = \sum_{o,\kappa,s,z,j} \frac{\tau_{do}^{\kappa,s,z,j} - 1}{\tau_{do}^{\kappa,s,z,j}} \cdot X_{do}^{\text{sea},\kappa,s,z,j\prime}$$

(ii) Global redistribution. Alternatively, if the policy generates benefits that accrue globally (e.g., emissions reductions from a carbon tax), revenue may be pooled and distributed by GDP share:

$$R_d^{\text{policy}} = \frac{w_d L_d}{\sum_c w_c L_c} \cdot R^{\text{policy}}.$$

Country d's income in the baseline is thus given by

$$I_d = w_d L_d + R_d^{\textbf{tariff}} + D_d$$

where R_d^{tariff} is standard tariff revenue from non-transport trade barriers. Country d's counterfactual income is

$$I_d' = w_d' L_d + R_d^{\text{tariff'}} + \underbrace{R_d^{\text{policy}}}_{\text{policy revenue}} + \underbrace{\Delta I_d^{\text{ship}}}_{\text{ship income change}} + D_d$$

5.5. Equilibrium

Given world labor endowments $(L_d)_d$, government policies (tariffs, trade costs and any port fees encoded in $\bar{\phi}_{od}^{\,j}$ and $\delta_{od}^{\text{sea},\kappa^j,s,z,j}$), and fixed ship capacities $(DWT^{\kappa,s,z,\text{supply}})$, an equilibrium is a set of wages $(w_d)_d$, sectoral prices $(P_d^j)_{d,j}$, bilateral expenditure flows $(X_{od}^j)_{o,d,j}$, and ship-service prices $(\Xi^{\kappa,s,z})_{\kappa,s,z}$ such that: (i) households and firms optimize, (ii) goods markets clear, (iii) ship-service markets clear and (iv) the trade-balance conditions $(D_d)_d$ are satisfied.

5.6. Solving for counterfactual equilibria in changes

Following the exact-hat algebra of Dekle, Eaton, and Kortum (2008), let $\hat{x} \equiv x'/x$ denote the counterfactual-to-baseline ratio. The transport block now delivers a set of new endogenous hat variables; the ship-service price changes and the composite transport cost changes.

Transport-cost hats. For *air* and *land* modes we assume no policy or technology changes in our counterfactuals, so that delivered costs are unchanged:

$$\hat{d}_{od}^{m,j} = 1, \qquad m \in \{\text{air}, \text{land}\}.$$

For the *sea* leg and vessel types $v \in \{L, S\}$,

$$\hat{d}_{od}^{\text{sea},\kappa^j,s,z,j} = \hat{\delta}_{od}^{\text{sea},\kappa^j,s,z,j} \, \hat{\Xi}^{\kappa,s,z}.$$

Let $\omega_{od}^{\kappa^j,s,z,j}$ denote baseline builder-country shares inside the lowest nest (vessel-class and size), $\omega_{od}^{\kappa^j,s,j}$ denote baseline vessel-size shares within vessel classes, and $\omega_{od}^{m,j}$ the baseline mode shares.

Using the CES nests,

$$\hat{d}_{od}^{\text{sea},\kappa^j,s,j} = \left[\sum_{z \in \mathcal{Z}} \omega_{od}^{\kappa^j,s,z,j} \left(\hat{d}_{od}^{\text{sea},\kappa^j,s,z,j} \right)^{1-\zeta} \right]^{1/(1-\zeta)}$$

$$\hat{d}_{od}^{\text{sea},j} = \left[\sum_{s \in \{L,S\}} \omega_{od}^{\kappa^j,s,j} \left(\hat{d}_{od}^{\text{sea},\kappa^j,s,j} \right)^{1-\rho} \right]^{1/(1-\rho)}$$

$$\hat{d}_{od}^{j} = \left[\sum_{m \in \mathcal{M}} \omega_{od}^{m,j} \left(\hat{d}_{od}^{m,j} \right)^{1-\eta} \right]^{1/(1-\eta)}$$

Trade cost changes follow as

$$\hat{\phi}_{od}^{j} = \begin{cases} \hat{\bar{\phi}}_{od}^{j} \hat{d}_{od}^{j} & \text{for } j \in \mathcal{J}^{\text{G}} \text{ (tradable merchandise)} \\ \hat{\bar{\phi}}_{od}^{j} & \text{for } j \in \mathcal{J}^{\text{S}} \text{ (tradable services)} \\ \hat{\bar{\phi}}_{od}^{j} \text{ if } o = d \text{ and } 1 \text{ if } o \neq d & \text{for } j \in \mathcal{J}^{\text{N}} \text{ (non-tradable sectors)} \end{cases}$$

Vessel-service market clearing. Given new trade flows, we obtain the change in demand for ship services as

$$\widehat{DWT}_{od}^{\text{sea},\kappa^j,s,z,j} = \frac{(\hat{d}^{sea,\kappa^j,s,z,j})^{-\zeta}}{\sum_{u \in \mathcal{Z}} \tilde{\omega}^{\text{sea},\kappa^j,s,u} (\hat{d}_{od}^{\text{sea},\kappa^j,s,u})^{-\zeta}} \frac{(\hat{d}^{sea,\kappa^j,s,j})^{-\rho}}{\sum_{r \in \mathcal{S}} \tilde{\omega}^{\text{sea},\kappa^j,r,j} (d_{od}^{\text{sea},\kappa^j,r,j})^{-\rho}} \frac{(d^{sea,j})^{-\eta}}{\sum_{\tilde{m} \in \mathcal{M}} \tilde{\omega}^{\tilde{m},j} (d_{od}^{\tilde{m},j})^{-\eta}} \frac{\hat{X}_{od}^{j}}{\hat{P}_{d}^{j}}$$

where $\tilde{\omega}^{\text{sea},\kappa^j,s,z} = \frac{(d^{sea,\kappa^j,s,z,j})^{-\zeta}}{\sum_{u \in \mathcal{Z}} (d_{od}^{\text{sea},\kappa^j,s,u})^{-\zeta}}$, $\tilde{\omega}^{\text{sea},\kappa^j,s,j} = \frac{(d^{sea,\kappa^j,s,j})^{-\rho}}{\sum_{r \in \mathcal{S}} (d_{od}^{\text{sea},\kappa^j,r,j})^{-\rho}}$, and $\tilde{\omega}^{\tilde{m},j} = \frac{(d^{m,j})^{-\eta}}{\sum_{\tilde{m} \in \mathcal{M}} (d_{od}^{\tilde{m},j})^{-\eta}}$

New total demand for ship services by type is

$$DWT_{od}^{\kappa,s,z\prime} = \sum_{j \in \mathcal{J}_{\kappa^j=\kappa}^{\text{G}}} DWT_{od}^{\text{sea},\kappa^j,s,z,j} \widehat{DWT}_{od}^{\text{sea},\kappa^j,s,z,j}$$

The market-clearing conditions in hats are

$$\sum_{od} DWT_{od}^{\kappa,s,z\prime} = DWT^{\kappa,s,z,supply}$$

and jointly determine the endogenous ship-service price hats $(\hat{\bar{\Xi}}^{\kappa,s,z})$.

The counterfactual changes in total and bilateral expenditure, prices, and wages, which closely follow Caliendo and Parro (2015), are described in Appendix A.2.

5.7. Calibration and Data

We calibrate the model using the 2025 release of the OECD Inter-Country Input-Output (ICIO) database for the reference year 2022. This provides the baseline N-country, J-sector network of global value added, gross output, and bilateral trade.

Elasticities. The model features nested substitution elasticities at the mode, vessel size, and builder level, in addition to standard trade elasticities. First, we set the sector-specific trade elasticities θ^j, which govern substitution across origin countries, using the product-level estimates from Fontagné, Guimbard, and Orefice (2022). Second, we set the elasticity of substitution across transport modes to $\eta = 3$, following estimates by Tolva (2024). Third, we utilize our novel estimates from Section 4 for the maritime sector. We set the substitution elasticity across vessel size classes to $\rho = 6.9$, consistent with our Wald-DiD estimate of the trade elasticity.[16] Finally, we set the substitution elasticity across builder countries to $\zeta = 7$, reflecting a degree of substitutability between shipyards that is close to the substitutability of vessel size (but still observes $\zeta > \rho$). As as robustness analysis, we also present results using a very high degree of substitutability ($\zeta = 26$). For our robustness analysis featuring endogenous 'congestion' cost for air and land transportation we use an elasticity of $\varphi = .06$, corresponding to the air transportation congestion elasticity estimate obtained by Tolva (2024).

Production and Consumption Parameters. We recover structural parameters directly from the input-output tables. Labor cost shares β_d^j and input-output coefficients $\gamma_d^{k,j}$ are calibrated to match the value-added and intermediate input shares in gross output for each sector and country in 2022. Consumption shares α_d^j are set to match sectoral shares in final demand. We close the baseline model by setting exogenous trade deficits D_d to match observed trade imbalances in 2022.

Baseline ship service market shares (mode $\pi_{od}^{m,j}$, size $\tilde{\pi}_{od}^{s,j}$, and builder $\tilde{\pi}_{od}^{z,s,j}$) are constructed based on transport mode shares from UN Comtrade and by combining aggregate ICIO trade flows with our granular bill of lading and AIS vessel data. Vessels of each class are categorized as small or large depending on their types and deadweight tonnage. For container ships, we use the same categorization as in Section 4; Post-Panamax and Neo-Panamax ships are classified as large, Panamax and other smaller vessel are classified as small. About 20% of container ships are large

---

16. Our estimated coefficient of 5.9 represents the elasticity of trade values ($\rho - 1$), implying $\rho = 6.9$.

according to this definition. For other vessel classes we choose a cutoff value for the deadweight tonnes such that the same share of vessels is classified as large. The corresponding cutoff values are listed in Table A2.

6. Counterfactual scenarios and results

The quantitative trade model developed above allows us to capture the general equilibrium implications of policies affecting the global shipping sector which have differential incidence across vessels. Here, we focus on two pertinent shocks, namely, (i) a new port fees on Chinese-built vessels entering US ports and; (ii) the extension of the EU Emission Trading System to maritime transport.

6.1. Port fees shock on Chinese vessels

In early 2025, US President Donald Trump proposed steep port fees on Chinese-built vessels docking at US ports, framing the measure as both a national security response to China's dominance in the maritime sector and a strategy to revitalize American shipbuilding. Set to be implemented by November 2026, this policy marks a further escalation in US-China trade tensions as it targets not only traded goods but also the ships that carry them.

Crucially, the effects of this policy extend beyond China and the US since they affect US imports from other countries relying on Chinese ships. Moreover, by changing the relative demand for ships built in China, the policy generates global spillovers through changes in the price of vessel services. Our model is designed to capture such new trade policies and their global spillovers through the shipping market as well as more standard general equilibrium effects working through the goods markets, in a tractable manner.

We model this port fees as a 10% tax on imports of goods arriving in the US on Chinese-built vessels.[17] The revenue from the fee is redistributed to US consumers. Table 3 summarizes the effects of this policy on the US economy. In our main simulation, which assumes a moderate degree of substitutability across ship building countries ($\zeta = 7$), imports and exports decline by about 0.2-0.3%, real output falls by 0.6%. Real wages drop by 0.3%, but due to the redistributed income from port fees, the fall in real consumption is lower. For a high degree of substitutability across builder countries, the trade and production effects are halved and the real wage effect is 30% smaller, whereas the drop in real consumption is very similar. Adding congestion in air and land transportation hardly changes the results.

---

17. See Table A3 for the distribution of shipping services across ship types and size class for vessels built in China and the rest of the world.

Table 3 — US port fees scenario: Effects on the US economy (changes in %)

	ζ	Real consumption	Real wage	Real output	Exports	Imports
USA	7	-0.06	-0.31	-0.62	-0.33	-0.22
USA	26	-0.06	-0.19	-0.35	-0.17	-0.10
USA	7	-0.06	-0.32	-0.65	-0.35	-0.23

Note: Scenarios. Row 1: main scenario. Row 2: large ζ. Row 3: congestion of air and land transport ($\varphi = .06$).

**Figure 10 — Changes in vessel service prices under new US port fees**
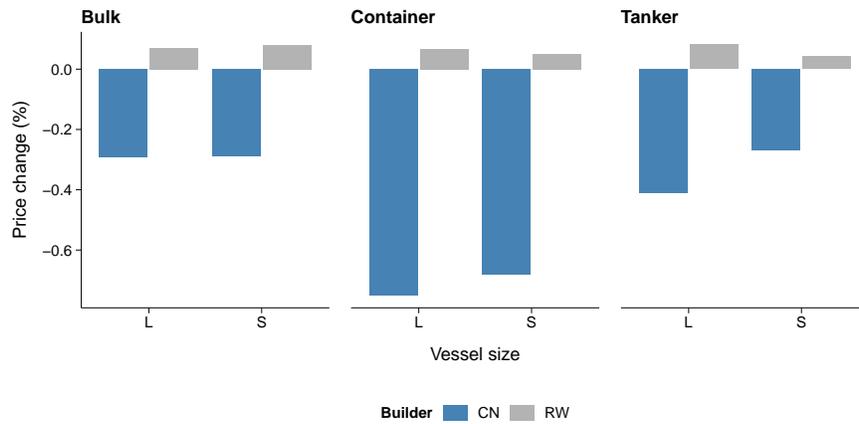
Note: $\zeta = 7$.

Besides the economic impact on the US, the port fee has global repercussions. These spillovers work through standard mechanisms in the goods market. US producers lose competitiveness due to higher prices of intermediate imports and the decline in economic activity reduces demand for foreign goods. Moreover, the port fees generates global spillovers through the shipping market. US importers switch away from Chinese-built vessels, towards vessels built elsewhere, and towards other modes of transportation. This reduces global demand for Chinese-built vessels, resulting in a decline in prices of these vessel services. Likewise, global demand for vessels built elsewhere increases, leading to a higher price. Figure 10 shows the price effects in the counterfactual equilibrium for the low builder-country substitution elasticity ($\zeta = 7$). Here, prices for the services of (large and small) vessels built in China drop by about 0.3%, .7%, .3-.4% for bulk ships, container ships, and tankers, respectively. Prices for services of alternative vessels increase, albeit by much less.

The decline in the price of Chinese-built vessels is advantageous for the rest of the world. Figure 11 shows the real consumption effects for the top-10 and bottom-10 countries in our model.

Among the most positively affected countries are the Southeast-Asian economies such as Myanmar, The Phillippines, and Thailand, gaining by more than 0.3% in terms of real consumption. Some South American and European countries gain substantially as well. Figure 12 shows that winning countries are those that rely most heavily on Chinese-built vessels in the baseline equilibrium, as measured by the share of production leaving those countries on Chinese-built vessels (panel a) or the share of expenditure arriving on Chinese-built vessels (panel b).

Figure 11 — Welfare changes: Top 10 and Bottom 10
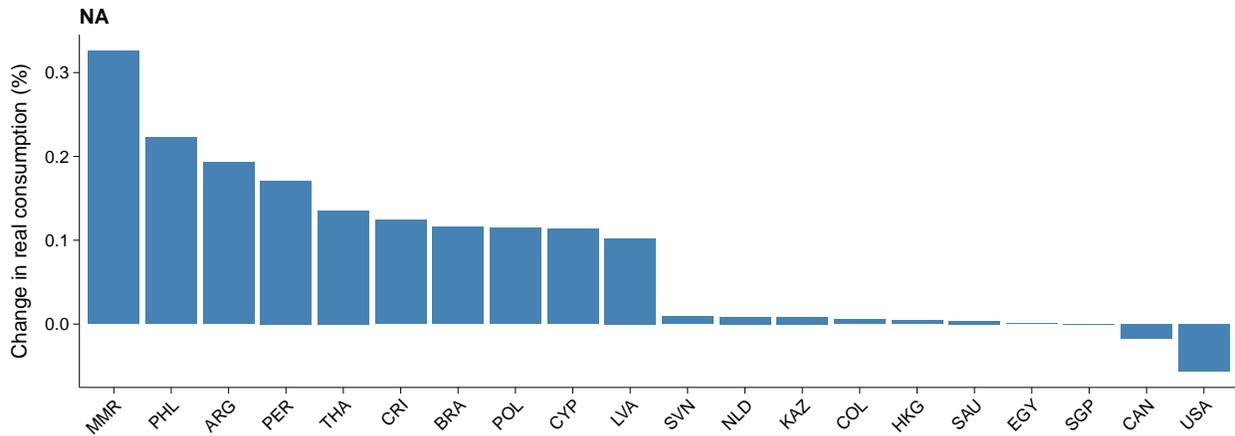

Note: $\zeta = 7$.

Figure 12 — Reliance on Chinese-built vessels in the baseline

(a) Export exposure
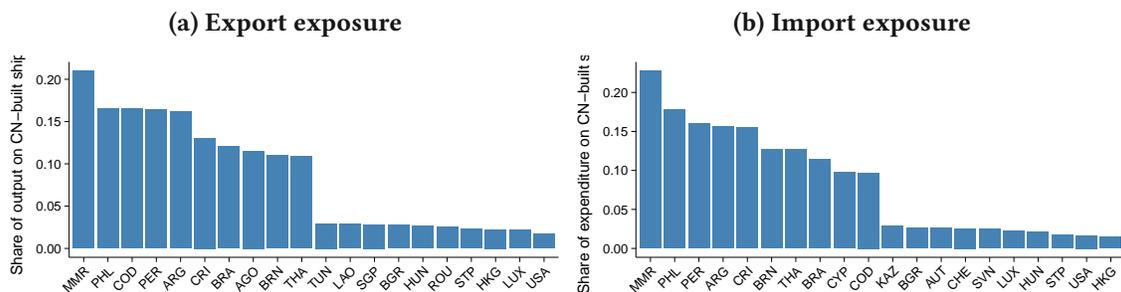
(b) Import exposure

Note: Export exposure is defined as the share of production leaving on Chinese-built vessels. Import exposure is defined as the share of expenditure arriving on Chinese-built vessels.

Next, we explore the heterogeneous effects of the policy on different sectors of the US economy. Due to China's dominance in the bulk shipping sector, US importers of commodities are particularly hard hit. Figure 13 (panel a) plots the share of US sectoral expenditure that arrives on Chinese-built vessels. Panel (b) contrasts this exposure with the price increases of sectoral goods in the counterfactual equilibrium. Prices of metals and minerals in the US increase the most, by

Figure 13 — Initial reliance on Chinese-built vessel by sector and corresponding price changes

(a) Import exposure baseline

(b) Counterfactual price change
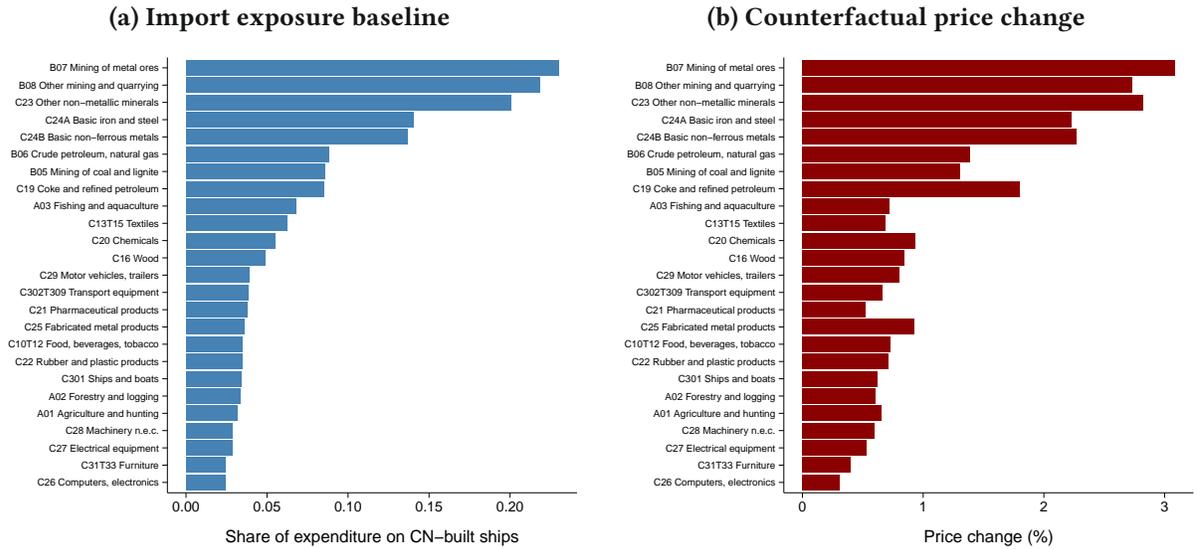
Note: Import exposure is defined as the share of expenditure arriving on Chinese-built vessels. $\zeta = 7$.

approximately 3%. Price of manufactured goods, for which the initial reliance on Chinese-built vessels is much lower, increase by less than 1%.

In sum, we find that US port fees on Chinese-built vessels harm consumers and producers in the US by increasing the prices of final and intermediate goods. The latter effect is concentrated in sectors that rely heavily on vessels that are predominantly built in China. Moreover, the US policy generates global spillovers through the market for shipping services. Prices for the services of Chinese-built vessels drop as US demand is diverted to ships from other builder countries. This benefits all other countries, in particular those that rely most heavily on Chinese-built vessels for their global imports and exports.

6.2. Pricing maritime emissions under the EU Emissions Trading System

In this counterfactual, we study the impact of the inclusion of maritime emissions in the EU Emissions Trading System (ETS). By 2027, vessel operators will be required to purchase allowances for 100% of their emissions related to voyages to or from ports in the European Economic Area (EEA) as well as voyages between EEA ports. The pricing of emissions will have heterogeneous effects across the fleets serving Europe, due to the variation in emission intensity across vessels.

Our model captures three relevant margins of heterogeneity: vessel type, vessel size, and builder country. We calculate the relative average emission intensity of the vessels across type-size-

builder country cells to inform the counterfactual cost changes in this scenario.[18]

The average emission intensity varies substantially across all three dimension in our model. For example, large container ships built in China emit 9% less than large container ships built elsewhere and 74% less than small container ships built in China. Our counterfactual cost shocks reflect these differences. We assume that the least emission-intensive type (large bulk ships, built either in China or elsewhere) experiences a cost increase of 2% due to emission pricing and then calibrate the cost shocks of other types according to their emission intensity relative to large bulk ships. On average across all products and partner countries of the EEA, this amounts to a 5% increase in trade costs on sea borne trade. Table A4 lists the cost changes (in ad-valorem equivalents) for sea-borne shipments to and from the European Economic Area (EEA) that we feed into the model.

In our main simulation, we assume that the revenue generated from the purchase of emission allowances accrues to the EEA country where targeted shipments arrive and depart. As an alternative, we assume that the revenue is redistributed to all countries, in proportion to their share in world GDP.[19]

Table 4 shows the economic effects of the policy on the EEA countries in our main scenario and under modified assumptions. Overall, the effects are negative and large. In the main scenario (row 1), imports, exports and real output of of EEA countries decrease by 5-6%. Real wages and real consumption drop by about 2.5% and 1%, respectively. Increasing the degree of substitutability of ships from different builder countries ameliorates the impacts (row 2), whereas allowing for the prices of air and land transportation to increase as demand moves away from sea transportation magnifies the negative effects (row 3). Redistributing the income of the policy globally rather than to EEA consumers leads to larger losses in real consumption, but to a smaller decline in real wages and real output.[20]

The cost of emission allowances reduces the demand for vessel services, in particular for small and more emission intensive vessel types. Figure 14 shows that trade on routes to and from the

---

18. To calculate the average emission intensity of vessel types we use the vessel-level emission data collected by the EU, available here: https://data.europa.eu/data/datasets/co2-emissions-data?locale=en. Emission intensity is defined as CO_2 emissions per transport work (g CO_2 / (tonnes carried * distance traveled).

19. Since our model captures neither current and future ownership structure of emission allowance, nor the global gains from emission reductions, nor the cost of emission abatement, it is not obvious how one should treat the revenue from emission allowances in our setting. The redistribution to EEA consumers describes a situation where EEA countries absorb the rents from higher prices of emission allowances and compensate consumers for higher cost of imports and exports caused by the scheme. The proportionate redistribution may best approximate a scenario where the revenue is spent on abatement. The increases in real consumption from this revenue may then be interpreted as the consumption equivalent welfare gain from lower global emissions.

20. This can be rationalized by de-facto redistribution of income from the EEA countries to the rest of the world, which depresses demand in the EEA, including imports, and spurs demand elsewhere, including demand for exports from the EEA. To ensure a fixed trade deficit, the wage in EEA countries is pushed up relative to the main scenario.

Table 4 — EU ETS scenario: Effects on the EEA economies (changes in %)

	ζ	Real consumption	Real wage	Real output	Exports	Imports
EEA	7	-0.93	-2.48	-5.30	-5.68	-5.61
EEA	26	-0.78	-2.94	-5.05	-4.74	-4.65
EEA	7	-0.96	-2.49	-5.32	-5.83	-5.76
EEA	7	-2.00	-1.00	-5.01	-5.44	-5.31

Note: Scenarios. Row 1: main scenario. Row 2: large ζ. Row 3: congestion of air and land transport ($\varphi = .06$). Row 4: global redistribution of income from emission allowances.

Figure 14 — Trade changes and dependence on small vessels
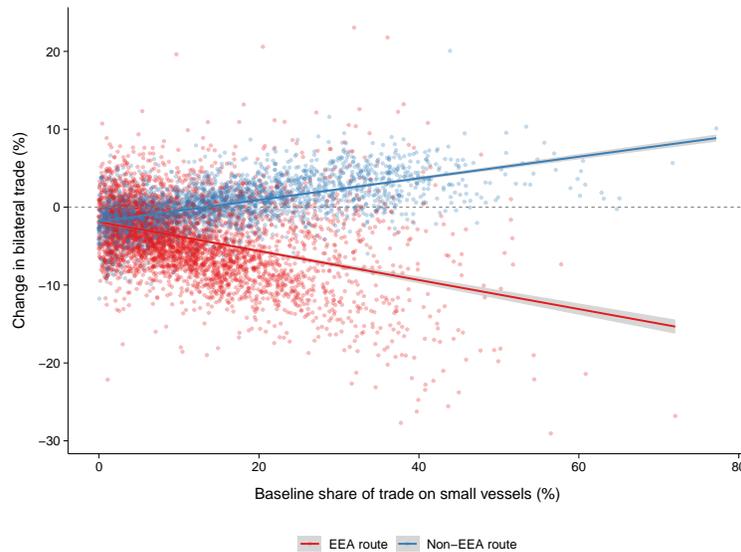
Note: $\zeta = 7$.

EEA tends to fall, in particular for routes that rely more strongly on small vessels in the baseline equilibrium (in red). As a consequence, the prices for services of small container ships and small tankers fall by about 4%. For other types of vessels, the price drop is only about 2%. The drop in the price of vessel services generates positive spillovers on third countries. Figure 14 shows that trade increases on routes connecting Non-EEA countries, and the more so the greater the initial reliance of these routes on small vessels (in blue).

Figure 15 — Changes in vessel service prices
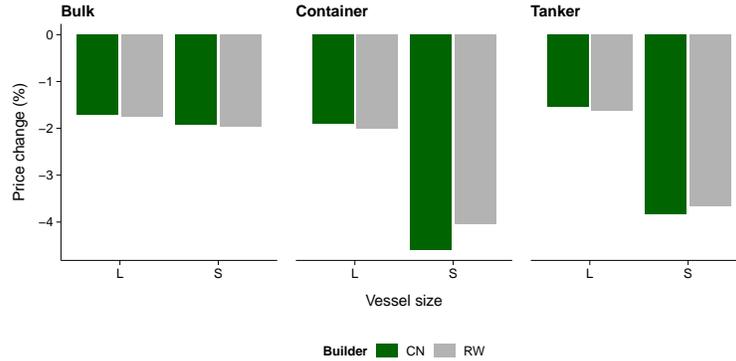
Note: $\zeta = 7$.

6.3. Discussion

Our counterfactual results show how policies targeting maritime transportation affect global trade and production. There are several points we would like to stress with regard to the interpretation of the quantitative results. First, in both scenarios, we have made an ad-hoc assumption about the level of the shocks because the ad-valorem equivalent of the port fees in the first scenario and the price of emissions as well as the share of emission-related cost in total vessel operating cost in the second scenario are unknown. Hence, the magnitude of our predicted changes is interpretable only in relation to the assumed shock size. Instead, the strength of our analysis lies in predicting the heterogeneity of effects across countries, sectors, or vessel types and the relative importance of novel mechanisms, such as global spillovers of policies through the market for different types of vessel services.

Second, both the policies we study aim to change economic outcomes in the long run, such as the distribution of ship-building activity across the world or the emission intensity of the global fleet. Due to the sluggishness of vessel building and vessel turnover, such effects will take decades to materialize. Our model takes the size and characteristics of the global vessel fleet as given. Hence, our results should be interpreted as median-run consequences of policies that target the long run.

Third, our EU ETS scenario focuses on the economic cost for EEA countries and leaves aside the environmental gains. Our negative real consumption effects should not be interpreted as evidence against the proclaimed social desirability of the scheme. However, we think that the two core novel mechanisms isolated by our model for the medium-run are relevant for the debate. As we show, emission-intensive ships are rerouted to non-EEA routes as a response to the policy. Moreover, the lower global prices of services provided by emission-intensive vessels partly compensate for the cost increases due to emission allowances also on EEA routes. Both mechanisms

limit the effectiveness of the scheme with regard to a reduction of global emissions, at least in the medium run.

7. Conclusion

The results in this paper highlight the importance of modeling shipping markets in richer detail when evaluating trade policies and transportation shocks. Measures that target specific vessel categories reshape the demand for vessel services and generate important global spillovers. By integrating endogenous vessel choice into a quantitative general equilibrium framework, we provide a tractable approach to studying these effects and their welfare consequences. Future work could extend this framework to dynamic fleet investment, further deepening our understanding of how the maritime transport system mediates the transmission of trade shocks in the global economy.

References

Alfaro, Laura, Harald Fadinger, Jan S Schymik, and Gede Virananda. 2025. *Trade and Industrial Policy in Supply Chains: Directed Technological Change in Rare Earths.* Technical report. National Bureau of Economic Research. (Cited on page 5).

Allen, Treb, and Costas Arkolakis. 2022. "The Welfare Effects of Transportation Infrastructure Improvements." *The Review of Economic Studies* 89, no. 6 (February): 2911–2957. ISSN: 0034-6527. https://doi.org/10.1093/restud/rdac001. eprint: https://academic.oup.com/restud/article-pdf/89/6/2911/46869678/rdac001.pdf. https://doi.org/10.1093/restud/rdac001. (Cited on pages 4, 25).

Ardelean, Adina, and Volodymyr Lugovskyy. 2023. "It pays to be big: Price discrimination in maritime shipping." *European Economic Review* 153:104403. (Cited on page 4).

Asturias, Jose. 2020. "Endogenous transportation costs." *European Economic Review* 123:103366. ISSN: 0014-2921. https://doi.org/https://doi.org/10.1016/j.euroecorev.2019.103366. https://www.sciencedirect.com/science/article/pii/S0014292119302272. (Cited on page 5).

Aytun, Uğur, Julian Hinz, and Cem Özgüzel. 2025. "Shooting down trade: Firm-level effects of embargoes." *Journal of Economic Behavior & Organization* 231:106821. (Cited on page 5).

Baldwin, Richard, Rikard Forslid, Philippe Martin, Gianmarco Ottaviano, and Frederic Robert-Nicoud. 2003. "10 The core–periphery model: key features and effects." In *The monopolistic competition revolution in retrospect,* 213. Cambridge University Press. (Cited on page 43).

Bonadio, Barthélémy. 2024. "Ports vs. roads: infrastructure, market access and regional outcomes," (cited on page 4).

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. "Geography, Transportation, and Endogenous Trade Costs." *Econometrica* 88 (2): 657–691. https://doi.org/https://doi.org/10.3982/ECTA15455. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15455. https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15455. (Cited on page 4).

———. 2024. *Investment in infrastructure and trade: The case of ports.* Technical report. National Bureau of Economic Research. (Cited on page 4).

Caliendo, Lorenzo, and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *Review of Economic Studies* 82 (1): 1–44. (Cited on pages 5, 21, 30).

Chowdhry, Sonali, Julian Hinz, Katrin Kamin, and Joschka Wanner. 2024. "Brothers in arms: The value of coalitions in sanctions regimes." *Economic Policy* 39 (118): 471–512. (Cited on page 5).

Clayton, Christopher, Antonio Coppola, Matteo Maggiori, and Jesse Schreger. 2025. "Geoeconomic Pressure." *Columbia Business School Research Paper Forthcoming, Stanford University Graduate School of Business Research Paper Forthcoming,* (cited on page 5).

Cristoforoni, Enrico, Marco Errico, Federico Rodari, and Edoardo Tolva. 2025. "Oligopolies in Trade and Transportation: Implications for the Gains from Trade," (cited on page 4).

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2018. "Fuzzy differences-in-differences." *The Review of Economic Studies* 85 (2): 999–1028. (Cited on page 4).

Dekle, Robert, Jonathan Eaton, and Samuel Kortum. 2008. "Global rebalancing with gravity: Measuring the burden of adjustment." *IMF Economic Review* 55 (3): 511–540. (Cited on page 28).

Do, Anh D, Sharat Ganapati, Woan Foong Wong, and Oren Ziv. 2025. *Transshipment hubs, trade, and supply chains.* Technical report. National Bureau of Economic Research. (Cited on page 4).

Drott, Constantin, Stefan Goldbach, and Volker Nitsch. 2024. "The effects of sanctions on Russian banks in TARGET2 transactions data." *Journal of Economic Behavior & Organization* 219:38–51. (Cited on page 5).

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, geography, and trade." *Econometrica* 70 (5): 1741–1779. (Cited on pages 21, 43).

Farrell, Henry, and Abraham L Newman. 2019. "Weaponized interdependence: How global economic networks shape state coercion." *International security* 44 (1): 42–79. (Cited on page 5).

Felbermayr, Gabriel, Aleksandra Kirilakha, Constantinos Syropoulos, Erdal Yalcin, and Yoto V Yotov. 2020. "The global sanctions data base." *European Economic Review* 129:103561. (Cited on page 5).

Feyrer, James. 2021. "Distance, trade, and income—The 1967 to 1975 closing of the Suez canal as a natural experiment." *Journal of Development Economics* 153:102708. (Cited on page 4).

Fontagné, Lionel, Houssein Guimbard, and Gianluca Orefice. 2022. "Tariff-based product-level trade elasticities." *Journal of International Economics* 137:103593. ISSN: 0022-1996. https://doi.org/https://doi.org/10.1016/j.jinteco.2022.103593. https://www.sciencedirect.com/science/article/pii/S0022199622000253. (Cited on page 30).

Fuchs, Simon, and Woan Foong Wong. 2024. *Multimodal Transport Networks.* CESifo Working Paper Series 11362. CESifo. https://doi.org/None. https://ideas.repec.org/p/ces/ceswps/_11362.html. (Cited on page 25).

Ganapati, Sharat, Woan Foong Wong, and Oren Ziv. 2024. "Entrepot: Hubs, scale, and trade costs." *American Economic Journal: Macroeconomics* 16 (4): 239–278. (Cited on pages 4, 5).

Gopinath, Gita, Pierre-Olivier Gourinchas, Andrea F Presbitero, and Petia Topalova. 2025. "Changing global linkages: A new Cold War?" *Journal of International Economics* 153:104042. (Cited on page 5).

Harrigan, James. 2010. "Airplanes and comparative advantage." *Journal of International Economics* 82 (2): 181–194. (Cited on page 4).

Heiland, Inga, Andreas Moxnes, Karen Helene Ulltveit-Moe, and Yuan Zi. 2025. "Trade from space: Shipping networks and the global implications of local shocks." *Review of Economics and Statistics,* 1–45. (Cited on pages 4, 14).

Hummels, David L, and Georg Schaur. 2013. "Time as a trade barrier." *American Economic Review* 103 (7): 2935–2959. (Cited on page 4).

Koenig, Pamina, Sandra Poncet, Mathieu Sanch-Maritan, Claude Duvallet, and Yoann Pigné. 2024. "Sold to China: Container traffic in the Port of Piraeus." *Review of International Economics* 32 (2): 510–544. (Cited on page 4).

Ludwig, Philipp. 2025. *Can Unilateral Policy Decarbonize Maritime Trade?* Technical report. CE-Sifo Working Paper. (Cited on page 5).

Tolva, Edoardo. 2024. "One way or another: Modes of transport and international trade," (cited on pages 4, 5, 23, 25, 30).

Verschuur, Jasper, Elco E Koks, and JW Hall. 2022. "Ports' criticality in international trade and global supply-chains." *Nature Communications* 13 (1): 4351. (Cited on page 4).

Wong, Woan Foong. 2022. "The Round Trip Effect: Endogenous Transport Costs and International Trade." *American Economic Journal: Applied Economics* 14, no. 4 (October): 127–66. https://doi.org/10.1257/app.20190721. https://www.aeaweb.org/articles?id=10.1257/app.20190721. (Cited on page 5).

A. Appendix

A.1. Details of the Caliendo and Parro (2015) model

There are N countries, indexed o and d, and J sectors, indexed j and k. Production uses labor as the sole factor, which is mobile across sectors but not across countries. All markets are perfectly competitive. Sectors are either wholly tradable or non-tradable.

There are L_d representative households in each country that maximize their utility by consuming final goods C_d^j in the familiar Cobb-Douglas form

$$u(C_d) = \prod_{j=1}^J C_d^{j\ \alpha_d^j} \quad \text{with} \quad \sum_{j=1}^J \alpha_d^j = 1.$$

where α_d^j is the constant consumption share on industries j's goods and where the respective Cobb-Douglas price index for consumption is denoted as P_d. Intermediate goods $\omega^j \in [0,1]$ are produced in each sector j using labor and *composite* intermediate goods from all sectors. Let $\beta_d^j \in [0,1]$ denote the cost share of labor and $\gamma_d^{k,j} \in [0,1]$ the input share of sector k in sector j's intermediate, such that

$$q_d^j(\omega^j) = z_d^j\left(\omega^j\right)\left[l_d^j\left(\omega^j\right)\right]^{\beta_d^j} \left[\prod_{k=1}^J m_d^{k,j}\left(\omega^j\right)^{\gamma_d^{k,j}}\right]^{1-\beta_d^j} \tag{A.1}$$

where $z_d^j(\omega^j)$ is the overall efficiency of a producer, $l_d^j(\omega^j)$ is labor input, and $m_d^{k,j}(\omega^j)$ represent the composite intermediate goods from sector k used to produce ω^j. With constant returns to scale and perfectly competitive markets, unit cost are

$$c_d^j = \frac{\Upsilon_d^j w_d^{\beta_d^j}}{z_d^j(\omega^j)}\left[\prod_{k=1}^J (P_d^k)^{\gamma_d^{k,j}}\right]^{1-\beta_d^j}$$

where P_d^k is the price of a composite intermediate good from sector k, and the constant $\Upsilon_d^j = (\beta_d^j)^{-\beta_d^j}\prod_{k=1}^J(\gamma_d^{k,j} - \beta_d^j\gamma_d^{k,j})^{-\gamma_d^{k,j}+\beta_d^j\gamma_d^{k,j}}$. Hence, the cost of the input bundle depends on wages and the prices of *all* composite intermediate goods in the economy. Producers of composite intermediate goods supply Q_d^j at minimum costs by purchasing intermediate goods ω^j from the lowest cost supplier across countries, so that

$$Q_d^j = \left[\int r_d^j(\omega^j)^{1-1/\sigma^j}d\omega^j\right]^{\sigma^j/(\sigma^j-1)}.$$

$\sigma^j > 0$ is the elasticity of substitution across intermediate goods within sector j, and $r_d^j(\omega^j)$ the demand for intermediate goods ω^j from the lowest cost supplier such that

$$r_d^j(\omega^j) = \left(\frac{p_d^j(\omega^j)}{P_d^j} \right)^{-\sigma^j} Q_d^j$$

where P_d^j is the unit price of the composite intermediate good

$$P_d^j = \left[\int p_d^j(\omega^j)^{1-\sigma^j} d\omega^j \right]^{1/(1-\sigma^j)}$$

and $p_d^j(\omega^j)$ denotes the lowest price of intermediate good ω^j in d across all possible origin locations, i.e.

$$p_d^j = \min_o \left\{ p_{od}^j \right\}. \tag{A.2}$$

Composite intermediate goods are used in the production of intermediate goods ω^j and as the final good in consumption as C_d^j, so that the market clearing condition is written as

$$Q_d^j = C_d^j + \sum_{k=1}^{J} \int m_d^{j,k}(\omega^j) d\omega^j \tag{A.3}$$

Trade in goods is costly, such that the offered price of ω^j from o in d is given by

$$p_{od}^j = \phi_{od}^j \cdot \frac{c_o^j}{z_o^j(\omega^j)} \tag{A.4}$$

where ϕ_{od}^j denote generic bilateral sector-specific trade frictions.[21]

Below we will derive the trade frictions for goods and commodities from a model of the multi-modal transportation sector.

Ricardian comparative advantage is induced à la Eaton and Kortum (2002) through a country-specific idiosyncratic productivity draw z^j from a Fréchet distribution.[22]

---

21. The "phiness" of trade à la Baldwin et al. 2003.

22. The productivity distribution is characterized by a location parameter λ_o^j that varies by country and sector inducing *absolute* advantage, and a shape parameter θ^j that varies by sector determining *comparative* advantage. θ^j describes the elasticity of trade to trade costs.

The price of the composite sector-j good in country d is then

$$P_d^j = A^j \left[\sum_{o=1}^{N} \lambda_o^j (c_o^j \phi_{od}^j)^{-\theta^j} \right]^{-1/\theta^j} \tag{A.5}$$

which, for a non-tradable or embargoed sector towards *all* non-domestic sources collapses to

$$P_d^j = A^j (\lambda_d^j)^{-1/\theta^j} c_d^j, \tag{A.6}$$

where $A^j = \Gamma(\xi^j)^{1/(1-\sigma^j)}$ with $\Gamma(\xi^j)$ being a Gamma function evaluated at $\xi^j = 1 + (1 - \sigma^j)/\theta^j$. Total expenditures on goods from sector j in country d are given by $X_d^j = P_d^j Q_d^j$. The expenditure on those goods originating from country o is called X_{od}^j, such that the share of j from o in d is $\pi_{od}^j = X_{od}^j / X_d^j$. This share can also be expressed as

$$\pi_{od}^j = \frac{\lambda_o^j (c_o^j \phi_{od}^j)^{-\theta^j}}{\sum_{h=1}^{N} \lambda_h^j (c_h^j \phi_{hd}^j)^{-\theta^j}}. \tag{A.7}$$

Total expenditures on goods from sector j are the sum of firms' and households' expenditures on the composite intermediate good, either as input to production or for final consumption:

$$X_d^j = \sum_{k=1}^{J} (1 - \beta_d^k) \gamma_d^{j,k} \sum_{o=1}^{N} X_o^k \frac{\pi_{do}^k}{\phi_{do}^k} + \alpha_d^j I_d. \tag{A.8}$$

Aggregate income in country d is

$$I_d = w_d L_d + R_d + D_d,$$

where $w_d L_d$ is labor income, R_d is net government revenue from trade policies (e.g., vessel fees) plus profits from the shipping sector, and D_d denotes the aggregate trade balance.

Sectoral trade balance is the difference between imports and exports,

$$D_d^j = \sum_{o=1}^{N} X_{od}^j - \sum_{o=1}^{N} X_{do}^j, \tag{A.9}$$

and aggregate trade balance is $D_d = \sum_{j=1}^{J} D_d^j$ with $\sum_{d=1}^{N} D_d = 0$. The vector $(D_d)_d$ is treated as exogenous, while sectoral trade balances $(D_d^j)_{d,j}$ are endogenously determined.

A.2. Solving for counterfactual equilibria in changes; additional details

The determination of counterfactual changes in prices, wages, and expenditures are unchanged relative to the standard EK system:

$$\text{Input costs} \qquad \hat{c}_d^j = \hat{w}_d^{\beta_d^j} \Big(\prod_{k=1}^J [\hat{P}_d^k]^{\gamma_d^{k,j}} \Big)^{1-\beta_d^j}$$

$$\text{Prices} \qquad \hat{P}_d^j = \Big(\sum_{o=1}^N \pi_{od}^j \, [\hat{c}_o^j \, \hat{\phi}_{od}^j]^{-\theta^j} \Big)^{-1/\theta^j}$$

$$\text{Trade shares} \qquad \pi_{od}^{j\prime} = \pi_{od}^j \Big(\frac{\hat{c}_o^j \, \hat{\phi}_{od}^j}{\hat{P}_d^j} \Big)^{-\theta^j}$$

$$\text{Expenditures} \qquad X_d^{j\prime} = \sum_{k=1}^J (1-\beta_d^k)\gamma_d^{j,k} \sum_{o=1}^N \frac{\pi_{do}^{k\prime}}{\phi_{do}^{k\prime}} X_o^{k\prime} + \alpha_d^j I_d'$$
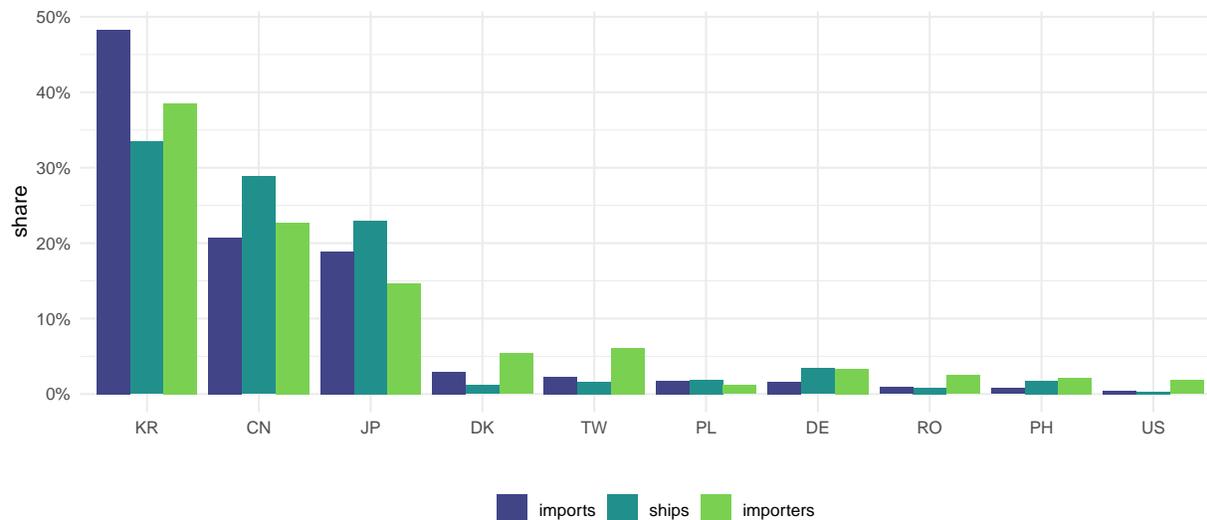
$$\text{Income} \qquad I_d' = \hat{w}_d \, w_d L_d + R_d' \; - \; D_d$$

$$\text{Trade balance} \qquad D_d = \sum_{j,o} \frac{\pi_{od}^{j\prime}}{\phi_{od}^{j\prime}} X_d^{j\prime} \; - \; \sum_{j,o} \frac{\pi_{do}^{j\prime}}{\phi_{do}^{j\prime}} X_o^{j\prime}$$

$$\text{Wage update} \qquad \hat{w}_o = \frac{1}{w_o L_o} \sum_{k=1}^J \Big(\beta_o^k \sum_d \frac{\pi_{od}^{k\prime}}{\phi_d^{k\prime}} X_d^{k\prime} \Big), \qquad \sum_d D_d = 0.$$

A.3. Additional results

Figure A1—Share of vessels in US maritime imports by builder country (2024)


Note: This figure shows the ten largest shipbuilding countries and their respective vessels' share in (i) US maritime import value (ii) ships arriving at US ports and; (iii) US importers receiving shipments by sea in 2024.

Table A1—Transition probabilities across size classes

type	stay	adjacent	jump
Bulker	0.656	0.330	0.014
Cargo	0.878	0.106	0.016
Tanker	0.828	0.108	0.065

Note: The table above reports the probabilities of remaining within the same vessel size class, switching to an adjacent size class, or making a larger jump across size classes when carriers replace a vessel between consecutive shipments for a given US importer–product–origin country triplet. Probabilities are computed separately for dry bulk carriers, tankers, and general cargo vessels, using US maritime import transactions from 2020–2024.

Figure A2 — Share of total port calls made by China-built vessels (2024)
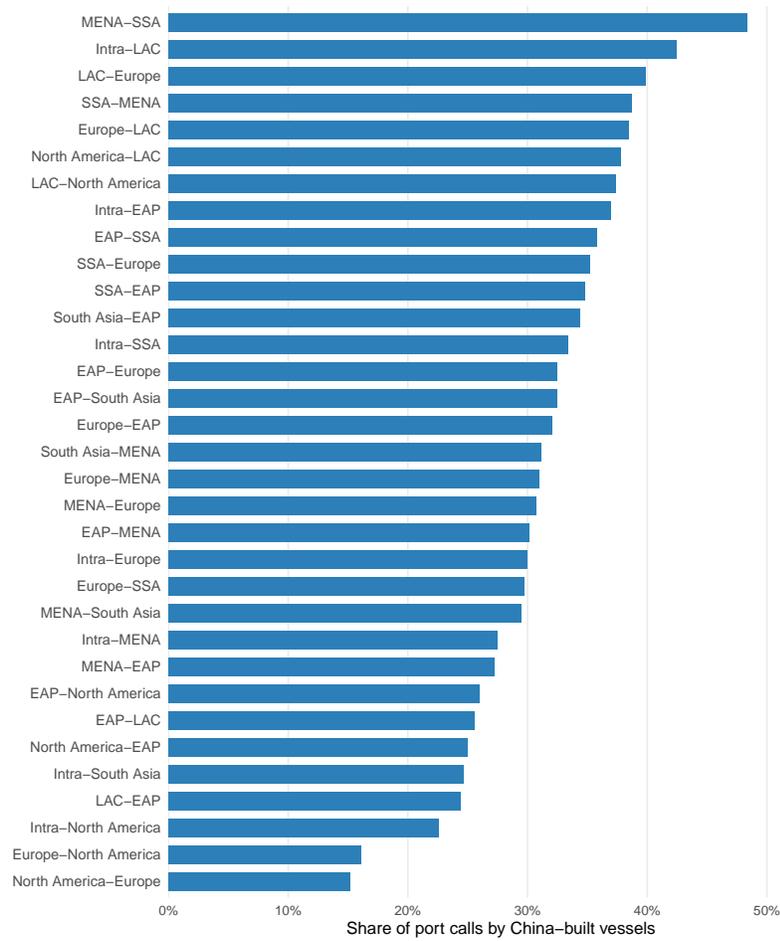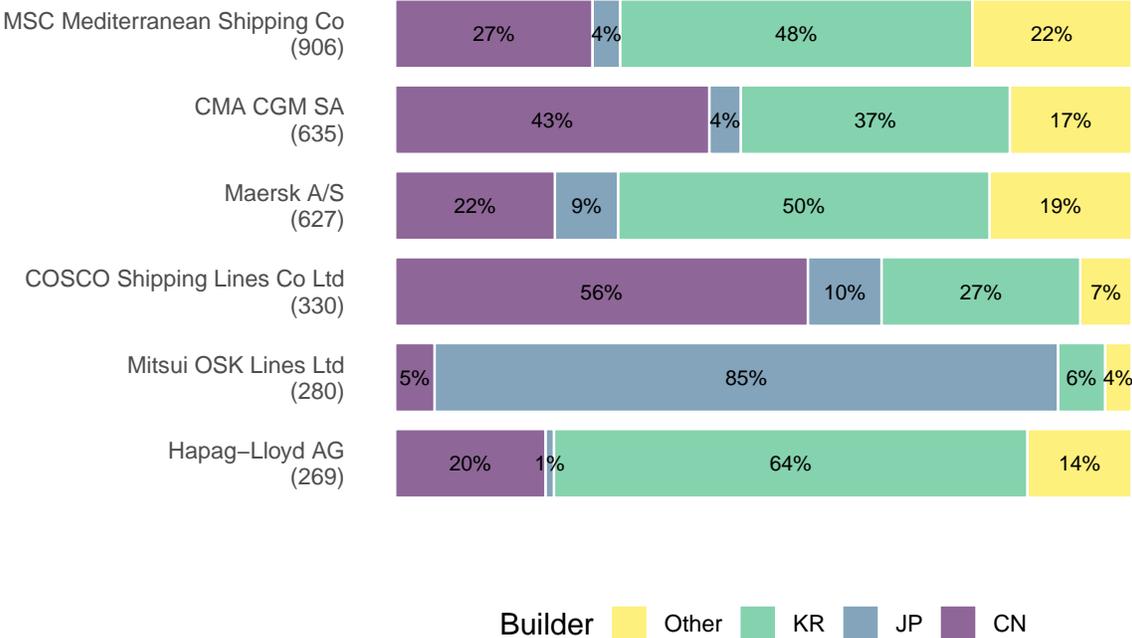

Note: This figure reports the share of port calls made by China-built container ships across various trade corridors in 2024. Each port call is assigned to a trade corridor based on the departure and arrival region. For corridors with at least 500 observed calls in 2024, we then compute the fraction that is executed by ships built in China.

Table A2 — Size cutoff for small and large vessels

Ship type	size cutoff (in deadweight tonnes)
Bulk	81243
Container	67515
Tanker	8133

Figure A3 — Builder shares in the global fleet of major carriers

Carrier	CN	JP	KR	Other
MSC Mediterranean Shipping Co (906)	27%	4%	48%	22%
CMA CGM SA (635)	43%	4%	37%	17%
Maersk A/S (627)	22%	9%	50%	19%
COSCO Shipping Lines Co Ltd (330)	56%	10%	27%	7%
Mitsui OSK Lines Ltd (280)	5%	85%	6%	4%
Hapag–Lloyd AG (269)	20%	1%	64%	14%

Builder: Other | KR | JP | CN

Note: This figure shows the composition of the fleets operated by the six largest global carriers. For each operator, it shows the total number of vessels they operate (built since 1995) and the shares that are constructed in China, Japan, South Korea, or other countries.

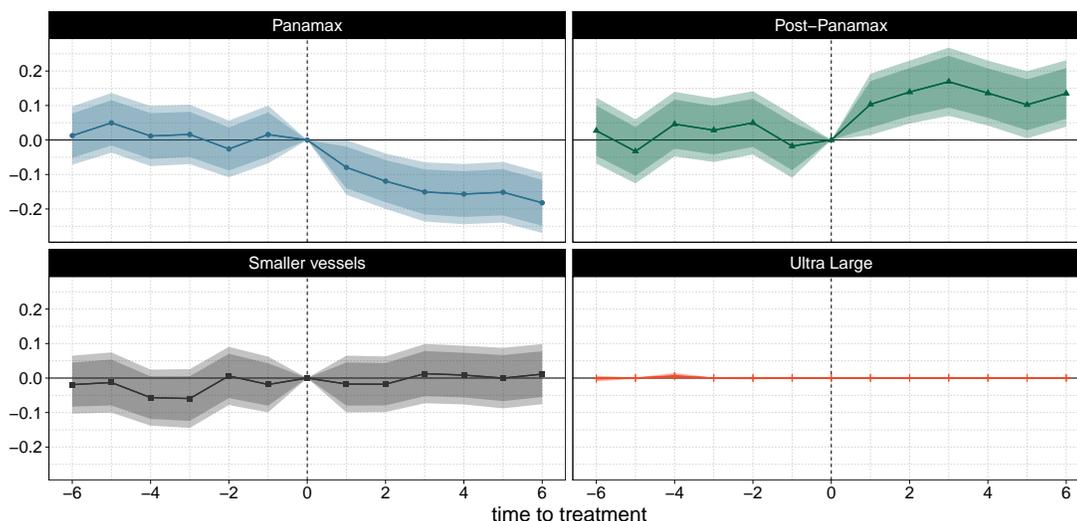## Figure A4 — Share of Neo-Panamax in US import volumes
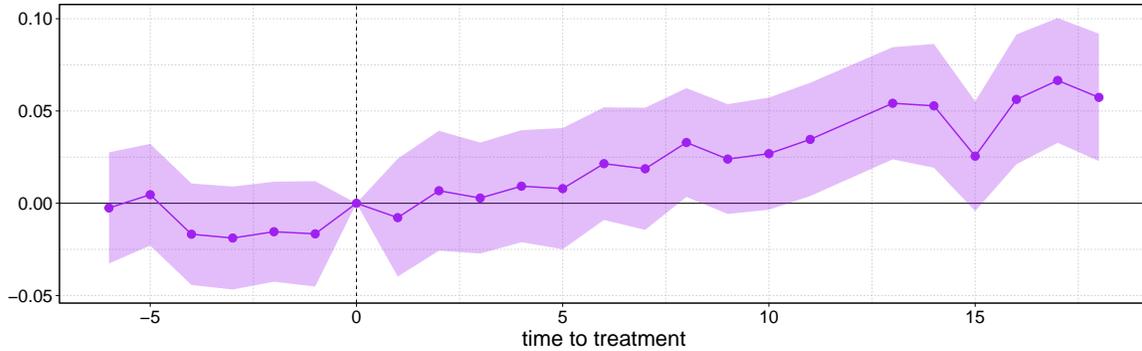
Note: This figure reports estimated β_k coefficients from equation (1) for months before and after the Panama Canal expansion in June 2016. The dependent variable corresponds to the share of US maritime import volume that is transported via Neo-Panamax container ships.

Figure A5 — Share of vessels in US import volumes, eligibility defined by beam
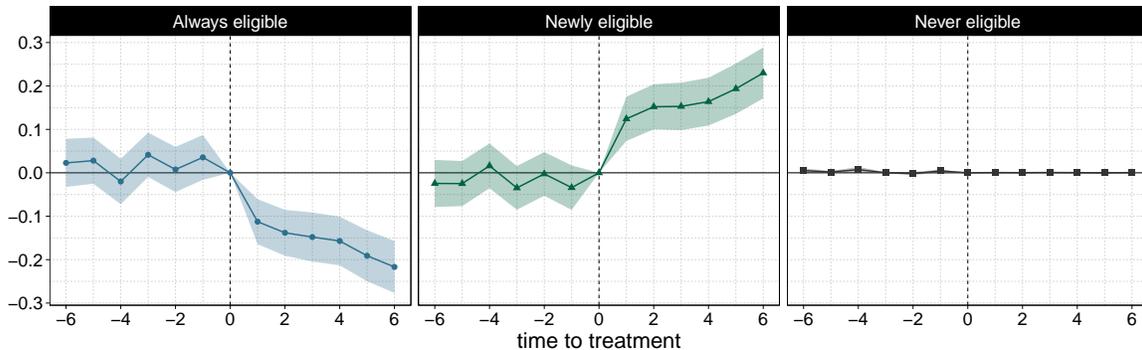
Note: This figure replicates the event study analysis in equation (1) using vessel beam rather than TEU capacity to define eligibility for the expanded Panama Canal locks. Vessels with beams below 32.3 meters were always eligible to transit, while those between 32.3 and 49 meters became newly eligible after the 2016 expansion. Vessels with beams beyond this threshold cannot be accommodated even after the Canal expansion.

Figure A6 — Share of vessels in US import values

Note: This figure reports estimated β_k coefficients from equation (1) for the months before and after the Panama Canal expansion in June 2016. In each panel, the dependent variable corresponds to the share of US maritime import values that is transported via vessels of the given size category.

Table A3 — Share of Chinese vessels in shipping services

Ship type	Size class	Builder country	DWT (billion)	Share (%)
Bulk	Large	CHN	34.2	44.9
Bulk	Large	ROW	42.0	55.1
Bulk	Small	CHN	50.2	45.6
Bulk	Small	ROW	60.0	54.4
Container	Large	CHN	149.5	15.7
Container	Large	ROW	801.7	84.3
Container	Small	CHN	191.4	32.1
Container	Small	ROW	405.5	67.9
Tanker	Large	CHN	16.3	20.4
Tanker	Large	ROW	63.4	79.6
Tanker	Small	CHN	14.2	22.5
Tanker	Small	ROW	48.8	77.5

Table A4 — Carbon tax scenario input (%)

Ship type	Size class	Builder country	Cost increase (%)
Bulk	Large	CHN	2
Bulk	Large	ROW	2
Container	Large	CHN	3
Container	Large	ROW	3.2
Tanker	Large	CHN	2.4
Tanker	Large	ROW	2.8
Bulk	Small	CHN	3
Bulk	Small	ROW	3
Container	Small	CHN	11.2
Container	Small	ROW	9.2
Tanker	Small	CHN	10
Tanker	Small	ROW	8.6