Can Unilateral Policy Decarbonize Maritime Trade?

Philipp Ludwig*

November 2024

Abstract

The international shipping sector is a vital part of the global trading system but also a large emitter of carbon dioxide emissions. In the absence of a multilateral carbon policy for the shipping sector, different countries are starting to impose unilateral measures to decarbonize maritime trade. This paper investigates the impact of unilateral policy on global carbon emissions and welfare by introducing heterogeneous transport technology to a quantitative model of trade. The framework emphasizes the role of transport providers which allocate clean and dirty vessels to shipping routes and thereby determine trade flows, transport costs and emissions in the entire network. Unilateral policy interferes with the allocation process by altering the cost of transport or restricting port access to specific vessel types. Using unique data on ship-level fuel consumption and network traffic, I study the impact of upcoming unilateral policy in the EU. I find that carbon taxes can achieve sizeable emission savings as long as the supply of transport services is sufficiently elastic. Emission savings, however, are not large enough to compensate for lower levels of trade, resulting in an aggregate loss of welfare.

Keywords: Carbon emissions, container shipping, transport network, unilateral policy
JEL Classification: F18, R12, R13, R41, H23, Q52, Q56

*I would like to thank Peter Egger, Farid Farrokhi, Teresa Fort, Richard Hornbeck, Jay Hyun, Nuno Limao, Kalina Manova, Isabelle Méjean, David Nagy, Ralph Ossa, Charly Porcher, Stephen Redding, Robert Staiger, David Torun, Ariel Weinberger, Zhihong Yu, and participants of the GEP/CEPR Postgraduate Conference, the CESifo Area Conference on the Global Economy, Rocky Mountain Empirical Trade Conference, LSE Environment Week and Trade Lunch Seminars at KU Leuven, Georgetown, Science Po, Paris 1 and ETH Zurich for valuable comments and suggestions. Special thanks go to my advisor Hylke Vandenbussche for her continued mentoring and encouragement, and to Sharat Ganapati for his hospitality and invaluable support during a research stay at Georgetown where this paper was developed. I gratefully acknowledge funding from the Research Foundation Flanders (FWO) and thank S&P Global for providing important data used in this project. Any remaining errors are my own.

*Department of Economics, KU Leuven. philipp.ludwig@kuleuven.be

1 Introduction

International trade is a major driver of anthropogenic climate change. Around one-third of trade-related carbon emissions are created outside the factory gate and instead are directly linked to the international transport of goods (Copeland et al., 2022). The lion's share of global transport activity is maritime. Container ships alone account for roughly 40% of global trade in terms of value and volume[1] and emit over 230 million tons of carbon dioxide (CO2) every year (IMO, 2020). This makes the shipping sector an important lever for cutting global CO2 emissions.

Concerted efforts to reduce maritime transport emissions have had limited success. While the International Maritime Organization (IMO), a multilateral agency tasked with monitoring the shipping sector, has communicated clear emission targets[2] to align the sector with the Paris Agreement of 2015, carbon emissions have continued to increase over the reference period (UNCTAD, 2022). This failure to cut carbon emissions is linked to a lack of enforceability at the multilateral stage[3] and has created a new trend in environmental policy: addressing global transport emissions unilaterally.

Starting in January 2025, the EU will impose a carbon tax and fuel standards on all ships operating in EU waters.[4] This represents the first binding carbon policy targeting the shipping industry to date and similar proposals have since been put forward in the US.[5] These unilateral attempts to decarbonize maritime trade raise two important questions: (1) Are unilateral policies a suitable tool to cut carbon emissions in maritime trade? (2) What are the welfare implications for countries around the world? This paper addresses these questions in four steps which range from an empirical investigation of proprietary

[1] Maritime shipping accounts for 70% and 80% of global trade in terms of value and volume (Brancaccio et al., 2020). Among all seaborne trade, goods carried on container ships account for 60% and 50% in terms of value and volume (Heiland et al., 2022).

[2] The 2023 IMO GHG Strategy calls for a 20-30% reduction of annual greenhouse gas emissions in international shipping by 2030 relative to the baseline year of 2008. CO2 emissions have increased by 10-15% over the same period (UNCTAD, 2022).

[3] IMO proposals typically have to be enacted by the flag state under which a ship is sailing. The IMO however has no formal power to sanction individual member states in case of non-compliance.

[4] The EU policy is discussed in detail in section 4.3.

[5] The *Clean Shipping Act* (H.R.8336) was first introduced to the US House of Representatives in 2022 and would impose increasingly strict fuel standards for all journeys involving a US port.

shipping data to counterfactual exercises using a quantitative trade model.

To start with, I present unique microdata to study the determinants of transport emissions in maritime trade. The data combines vessel movements with hourly fuel consumption estimates of every container ship which allows me to track carbon emissions across the entire transport network. The descriptive analysis of the data bears two important insights: Fuel efficiency, measured in tons of fuel per container-kilometre, varies significantly within the fleet and is skewed towards traffic-intensive parts of the network. Transport providers therefore play an important role for global transport emissions by choosing how to allocate clean and dirty vessels across the network.

My first contribution is to establish a set of stylized facts that link the observed allocation behaviour of shippers to cost incentives created by a heterogeneous stock of transport technology. Ship-level fuel efficiency is the result of two underlying mechanisms: heterogeneity *between* ships, which is largely related to ship size, and heterogeneity *within* ships, determined by the share of transport capacity in use, also known as the load factor. Large ships in general consume less fuel per container-km, making them the most cost-effective and cleanest form of maritime transport. This advantage however vanishes if a ship is not operating at full capacity. As only certain segments of the global network experience sufficient transport demand to sustain high load factors, vessel operators are incentivized to allocate their cleanest ships to traffic-intensive parts of the network. To gauge the environmental impact of unilateral policy, we therefore need to account for the endogenous reallocation of transport technology and ensuing change in fuel efficiency across network segments.

I formalize this mechanism by integrating transport emissions into a structural model in which trade is routed through an endogenous transport network. I follow Ganapati et al. (2024) and combine a canonical trade model featuring Ricardian selection and input-output linkages (Caliendo and Parro, 2015) with a route selection model of Allen and Arkolakis (2022). In the model, trade flows are routed indirectly through a global transport network, mimicking the flow of containerized trade observed in the data.

The key novelty and second contribution of this paper is to introduce transport emis-

sions as the result of a technology allocation process in the transport sector. Transport services are provided by a heterogeneous fleet of vessels that travels along a fixed set of circular shipping routes. Shippers allocate vessels to routes to serve transport demand at the lowest possible cost, taking into account that fuel efficiency differs across vessels and varies with the load factor on each segment. These cost incentives result in an optimal allocation in which larger and more efficient vessels are allocated to network segments where transport demand is high, just as observed in the data. This generates a cost advantage that is passed on to consumers in the form of lower transport cost on these segments. As the allocation changes endogenously with the level of transport demand, so does transport cost and the carbon emissions created in the process. Expressing transport cost as the outcome of an endogenous allocation process has two important benefits: First, it explains why transport costs fall at higher levels of traffic and thereby provides a new microfoundation for scale effects in the maritime transport sector. Scale effects play an important role for the welfare effects of policy interventions (Asturias, 2020; Ganapati et al., 2024) and emerge naturally in my framework. At the same time, the allocation mechanism determines the fuel efficiency of transport services throughout the network which allows me to study the environmental consequences of carbon policy.

The theoretical framework allows me to study the environmental impact of two distinct types of policy measures. Cost-based interventions such as carbon taxes aim to lower transport emissions by reducing transport demand. I show that the efficacy of this approach is closely linked to the supply elasticity of transport capacity. If demand for transport services falls while the supply of vessels is inelastic, shippers are forced to deploy more ships than necessary which decreases load factors and can offset any emission savings of the policy. Overcapacity thus lowers the efficiency of transport services and reduces the environmental gains of unilateral policy interventions. Allocation-based interventions instead shape carbon emissions by imposing unilateral emission thresholds or fuel standards for all vessels operating in their jurisdiction. If this unilateral interference runs counter to the cost incentives of shippers, it causes an allocation distortion which lowers emissions locally but raises transport emissions globally. Unilateral interventions

in theory therefore do not guarantee any emission savings, but instead create complex trade-offs that can only be resolved empirically - an insight which marks the third contribution of this paper.

To study the impact of unilateral policy on global carbon emissions and welfare, I calibrate the model to the shipping network in 2018. I then perform counterfactual exercises based on unilateral carbon taxes and fuel standards proposed by the EU which yield three important results: (1) Accounting for the endogenous response of transport providers lowers emission savings by 50%. Ignoring changes in network efficiency due to a reallocation of vessels or falling load factors leads to a severe overestimation of environmental gains of the policy. (2) Carbon taxes achieve sizeable emission savings as long as shippers can avoid overcapacity by withdrawing vessels from their fleet. In the absence of overcapacity, a 10% EU carbon tax reduces maritime transport emissions by 7.8% and production emissions by 1.4%, saving a total of 22 million tons of CO2. While this points to important complementarities between both emission types, EU transport taxes shift global production to non-EU countries which raises the emission intensity per unit of output. Unilateral policy might therefore be more effective when initiated by emission-intensive, non-EU countries. (3) Emission savings are not large enough to raise welfare. I find that a 10% EU carbon tax decreases global income levels by 0.3%. This income effect exceeds the aggregate social benefit of emission savings by several orders of magnitude and results in a welfare loss across countries. Acting unilaterally can therefore push the sector towards decarbonization but only at an enormous cost to welfare.

Related literature: By combining detailed microdata with a structural model of trade, transport and emissions, this paper contributes to three broad strands of literature. A first strand is a large literature studying the interaction between trade and the environment. Research in this area has studied the determinants of carbon emissions (Copeland and Taylor, 1994; Antweiler et al., 2001), the impact of trade policy (Cherniwchan, 2017; Shapiro and Walker, 2018; Shapiro, 2020; Bombardini and Li, 2020), carbon leakage (Hanna, 2010; Chung, 2014; Aichele and Felbermayr, 2015; Fowlie and Reguant, 2018) and the optimal design of carbon policy (Nordhaus, 2015; Kortum and Weisbach, 2022; Farrokhi

and Lashkaripour, 2023). The vast majority of the literature however exclusively focuses on production emissions, ignoring the environmental externality created by the international transport of goods. By shedding light on transport emissions, this paper not only fills an important gap in the literature, but also allows us to study the joint response of transport and production emissions to changes in unilateral policy.

This paper also contributes to a literature studying trade networks. Acknowledging the indirect nature of trade has yielded important insights about the role of infrastructure investments (Redding and Turner, 2015; Heiland et al., 2022; Ducruet et al., 2024), modal choices (Coşar and Demir, 2018; Fuchs and Wong, 2022) and endogenous transport cost (Brancaccio et al., 2020; Fajgelbaum and Schaal, 2020; Ganapati et al., 2024; Wong, 2022; Allen and Arkolakis, 2022) which all shape how goods flow through the network and have important consequences for the propagation of shocks and the distribution of welfare. This paper offers a parsimonious way to introduce carbon emissions to a quantitative model of trade, allowing us to study the role of network features for environmental outcomes without losing the rich structure of theoretical frameworks developed in this literature.

The literature closest to this paper directly focuses on transport emissions. This burgeoning literature has generated valuable insights about transport emissions across different modes of homogenous transport technology (Cristea et al., 2013; Shapiro, 2016), emphasized health benefits of restricting local pollution near coastal settlements (Hansen-Lewis and Marcus, 2022) and warned against inter-modal leakage as a consequence of tightened environmental regulation (Lugovskyy et al., 2024). My setting differs from these papers in two important aspects. First, I limit attention to a single mode of transport but allow intra-modal technology to be highly heterogeneous. This reveals a novel form of carbon leakage that stems from a reallocation of transport technology within the same sector. Second, in contrast to local pollutants such as particulate matter or sulphur, carbon emissions are a global externality that affects everyone irrespective of the origin of pollution. This raises novel questions about the design of unilateral transport policy whose immediate impact is local but whose efficacy depends on the response of economic agents around the globe.

2 Data

This section documents the key determinants of global transport emissions in containerized trade. I start by presenting two proprietary datasets[6] which capture ship movements and transport emissions of the global fleet. Together they provide a complete picture of how goods travel through the global transport network and the environmental externality created in the process. I then use this rich source of information to study the key drivers of transport emissions in maritime shipping.

The first dataset contains port of call information that is recorded by Automatic Identification System (AIS) transceivers installed on container ships. AIS data is available for every container ship in the global fleet and records ship movements between 2014 and 2018. Each time a ship enters or leaves a port, a satellite records the current time, location and submersion of the vessel. This detailed information allows me to reconstruct every ship journey taken between a pair of ports, infer the volume of transported goods[7] and capture the entire flow of containerized trade in the transport network.

The second dataset contains hourly estimates of ship-level fuel consumption in 2018. These estimates are based on high-resolution AIS data and a detailed set of ship characteristics. I use these hourly estimates to measure the fuel efficiency of every container ship in the global fleet. Fuel efficiency is hereby defined as the amount of fuel a ship needs to consume to transport a single twenty-foot equivalent container (TEU) over the distance of one kilometre. This efficiency measure varies between ships and changes with the current load factor of each vessel as explained in detail in section 3.

Combining both datasets then allows me to observe the transport emissions associated with every individual trip taken by a container ship in the global transport network. To do so, I compute the total fuel consumption of each journey and use a constant emission factor published by the IMO (2020)[8] to translate fuel consumption to CO2. I therefore

[6] Both datasets have been provided to me by S&P Global and are described in detail in appendix A.1.

[7] Transport volume is calculated by comparing the ship's current submersion, as recorded by the satellite, to the ship's submersion when travelling empty, which is reported in a separate dataset capturing ship characteristics. The exact engineering formula is described in appendix A.1.

[8] I use a benchmark emission factor of 3.114 $gCO_2/gfuel$ as reported for low sulphur heavy fuel oil (LSHFO) in table 21 of the IMO's Fourth Greenhouse Gas Study 2020.

abstract from variation in fuel types when computing transport emissions.[9]

Aggregate transport emissions can be decomposed into two key determinants: (a) the average fuel efficiency of the global fleet and (b) the total level of transport demand measured in container-kilometres. Expressing transport demand in units of TEU-km is important when taking the perspective of transport providers because it reflects the amount of transport capacity required for a given service. Two shipments which involve the same number of containers but are transported over an uneven distance for example require different levels of capacity and will create different costs for transport providers.[10] The shipping industry fittingly refers to TEU-km as the level of 'transport work' which is the main measure of transport demand in this paper.

Figure 11 in appendix B.1 shows how both determinants have evolved between 2016-2018.[11] While average fleet efficiency has improved over time, aggregate transport demand has increased disproportionately during the same period, resulting in a rise in global transport emissions.[12] Rejuvenating the global fleet by the introduction of newer and cleaner vessels evidently has not been sufficient to reduce carbon emissions in the sector.

Can unilateral policy revert this trend? The answer to this question depends on how the two emission determinants, average fuel efficiency and transport demand, respond to policy interventions. In this paper, I focus on the short to medium run in which unilateral policy is unlikely to trigger the mass-rollout of a carbon-neutral propulsion technology.[13] I therefore assume that the stock of available technology remains largely unchanged[14] but still allow policy interventions to affect transport emissions via transport demand.

[9]The emission intensity of the two main fuel types, heavy fuel oil and liquefied natural gas, differs by 13% according to emission factors published in EU regulation 2023/1805. Differences in emission factors across marine fuel types are therefore small when compared to the large variation in fuel efficiency between and within vessels reported in section 3.

[10]Transporting 5,000 TEUs between New York and Boston requires a lot less transport work than shipping the same number of containers from New York to Rotterdam.

[11]While AIS data records all ship movements between 2014-2018, draught data which is required to compute cargo loads is only available from 2016.

[12]This pattern is not limited to the sample period. Transport emissions from containerized trade have increased consistently over the past decade as shown in table 4.31 of UNCTAD (2022).

[13]The average time to build a new container ship is three years. While several companies have started to test container ships using carbon-neutral propulsion technologies, these pilot programs are unlikely to achieve large-scale adoption in time for the IMO's 2030 emission target.

[14]I relax this assumption in section E.3 by increasing the share of clean ships in the global fleet.

A key insight of this paper is that the second emission determinant, fuel efficiency, continues to respond to policy interventions even if the number of clean and dirty vessels remains completely unchanged. To illustrate this point, we first need to understand the origins of vessel heterogeneity in the global fleet.

3 Stylized Facts

In this section, I present three stylized facts to illustrate how vessel heterogeneity shapes the allocation of technology in the global transport network. Large ships enjoy an efficiency advantage over smaller ships and all vessels become more efficient when operating at higher load factors. As load factors in the network increase with segment traffic, operators minimize fuel consumption by allocating their largest vessels to segments where transport demand is high. The allocation of technology in the network is therefore determined by the cost-saving incentives generated by heterogeneity in the global fleet and the distribution of traffic across network linkages. Allowing shippers to follow these incentives is crucial to minimize fuel consumption and global transport emissions. I then discuss the consequences of this allocation behaviour for the theoretical framework.

3.1 Transport technology is highly heterogeneous

Figure 1 shows the fuel efficiency of all 5,100 container ships active in 2018. Each panel illustrates an important dimension of heterogeneity. Panel 1a focuses on differences between ships and shows that (inverse) efficiency is increasing with ship size. A medium-sized vessel with a capacity of 5,000 TEU on average consumes 57% more fuel per container-km than a ship with a transport capacity of 20,000 TEU. Big ships can spread their fuel consumption over a larger number of container slots and on average achieve higher levels of fuel efficiency in the process. This cost advantage of large ships however is contingent on the load factor at which they operate. Fuel efficiency within ships falls dramatically if container slots remain empty. A 5,000 TEU vessel at full load can therefore be as efficient as the largest vessel in the fleet if the load factor of the latter is low. Panel 1b illustrates

Figure 1: Fuel efficiency of the global fleet of container ships

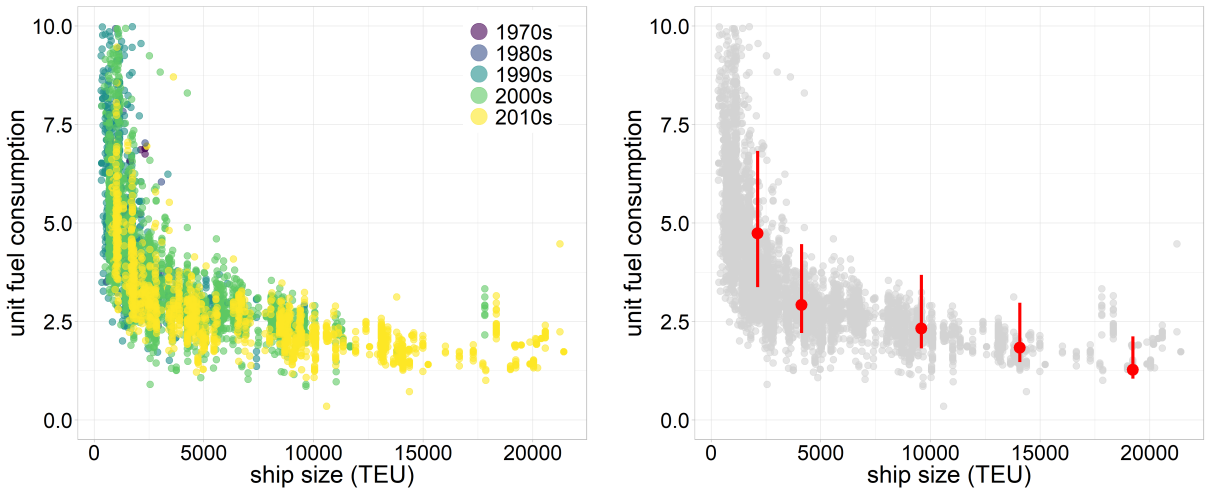

(a) between-ship heterogeneity

(b) within-ship heterogeneity

Note: Figure 1 shows the fuel efficiency of all active container ships in 2018. Efficiency is measured as the amount of fuel a ship needs to consume to transport a single container for one kilometre at sea ($= gfuel/(teu * km)$) and is based on data described in appendix A.1. Ship size is measured as TEU capacity. Colours in panel 1a represent the decade in which the ship was built. Bars in panel 1b show the possible range of fuel efficiency for selected vessels using within-ship load factor elasticity β estimated in Appendix D.1.

this pattern by plotting the efficiency range for select vessels as red bars[15] where the lower (upper) end of each bar refers to the ship's efficiency at full (zero) load. I summarize these observations as follows:

Stylized fact 1: *Fuel efficiency is highly heterogeneous between and within ships and increases with ship size and the load factor of each vessel.*

3.2 Load factors increase with leg activity

High load factors are crucial to keep the cost of transport services as low as possible. In practice, however, load factors vary significantly due to an uneven distribution of transport demand across the network. Figure 2 shows the relationship between transport demand and average load factors across different network segments. The point estimates of binscatter regressions across ship size bins indicate a strong log-linear relationship between load factors and segment traffic.[16] Ships operating on traffic-intensive parts

[15] The indicated range is based on the within-ship load factor elasticity estimated in appendix D.1.

[16] Segment traffic (or activity) here is measured in units of transport work. Traffic intensity therefore refers to the supplied amount of transport capacity on the segment. Supplied capacity increases if the number of transported containers or the port-to-port distance is large.

Figure 2: Load factors and transport work across network segments

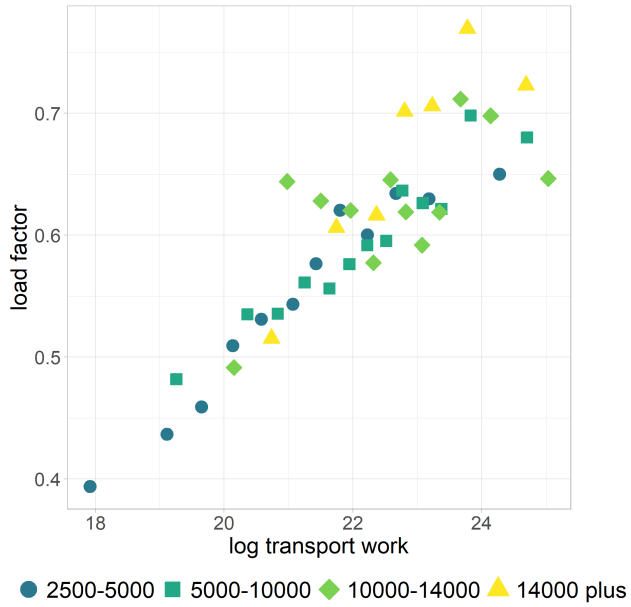

Note: Figure 2 plots point estimates of a binscatter regression following Cattaneo et al. (2023). Load factors and transport work have been aggregated to the segment level as described in appendix A. Separate graphs for each size class are available in appendix B.2. Ships are grouped into different bins according to their transport capacity in TEU, represented by different colours and shapes.

of the network on average experience higher load factors - a pattern that holds across ship types as shown in figure 12 in appendix B. This systematic relationship[17] results in efficiency gains for *any* ship operating on these segments. To summarize:

Stylized fact 2: *Container ships operating on traffic-intensive network segments on average experience higher load factors and fuel efficiency.*

3.3 Clean technology is allocated to traffic-intensive legs

Stylized facts 1 and 2 create a clear incentive structure for the allocation of ships in the maritime transport network. To exploit the efficiency differences in their fleet and save cost, vessel operators should allocate their largest ships to traffic-intensive network segments where load factors are high.

[17]The theoretical framework presented in section 4.2 rationalizes this positive relationship as the result of demand dispersion within observed shipping routes. The amount of transport capacity supplied to a route is determined by the route-segment with the largest transport demand. On the top route-segment this leads to a load factor of 1 while all other route-segments experience load factors less than 1. Network segments with high amounts of traffic on average also experience higher levels of demand within routes, making them more likely to be the top segment. This creates a positive correlation between segment demand and load factor.

Figure 3: Allocation of transport technology across network segments

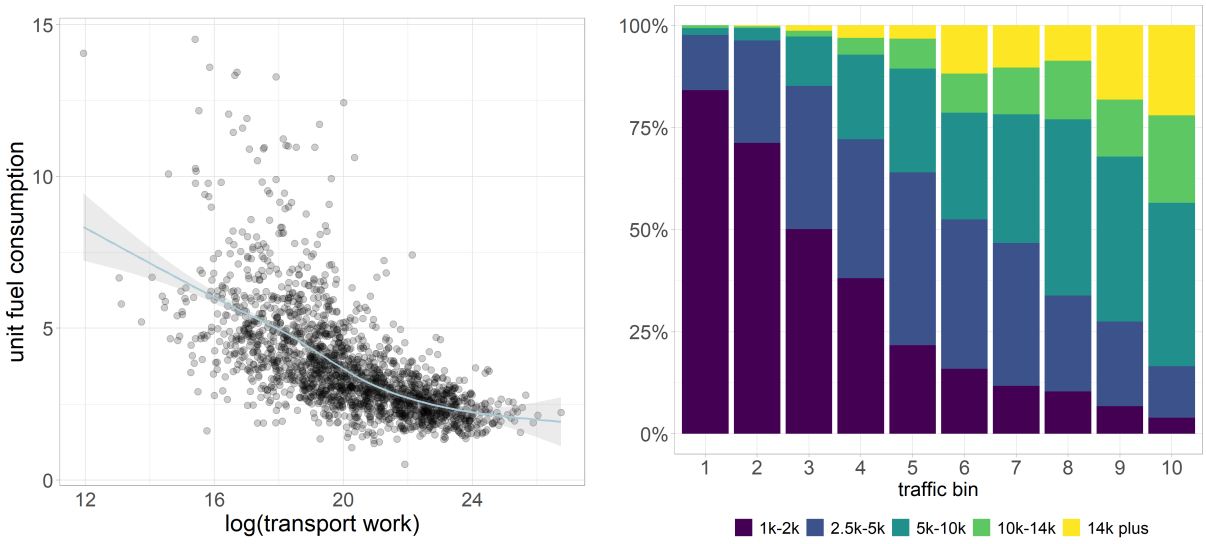

(a) fuel efficiency on the segment

(b) transport work share within segments

Note: Figure 6 depicts the observed allocation of container ships across network segments. Panel 6a shows the average fuel efficiency on network segments (y-axis) by segment traffic measured in units of transport work (x-axis). Panel 6b groups network segments into traffic bins ranging from low (bin = 1) to high (bin = 10). Colours refer to ship size bins and indicate what share of traffic within a traffic bin is served by vessels of a given size. Vessel size is hereby measured as total TEU capacity.

Figure 6 demonstrates that the observed allocation in the data is fully consistent with these incentives. Average fuel efficiency at the segment level systematically decreases with segment traffic (6a), precisely because big vessels (yellow and green bars) account for a larger share of transport work when comparing low to high traffic bins in the network (6b). In summary:

Stylized fact 3: *Large ships are allocated to traffic-intensive network segments - an allocation fully consistent with cost-saving incentives generated by stylized facts 1 and 2.*

Establishing a direct link between the observed allocation behaviour and cost-saving incentives of vessel operators generates two important insights. First, it provides an empirical underpinning for the theoretical framework presented in section 4. Here, I illustrate how the observed allocation of container ships can be replicated from two simple ingredients: a heterogeneous stock of technology at the ship-level and a fixed set of circular routes on which they operate. Second, it suggests that there exists an optimal allocation under which operators achieve the lowest possible cost by fully exploiting the cost-incentives laid out in stylized facts 1 and 2. Any alternative allocation which deviates from this bench-

mark raises aggregate fuel expenditure and transport emissions. The *allocative efficiency* of the global fleet can therefore be defined by the degree of control operators enjoy over the allocation of their vessels.

In practice, the allocative efficiency is constrained by infrastructure, policy, and other factors which restrict to which parts of the network ships can be allocated. A 20,000 TEU vessel will for example never be able to dock at a port which cannot accommodate ships of that size, even if doing so would be beneficial from the perspective of vessel operators. The allocation observed in figure 1a is therefore not unconditionally optimal, but instead reflects the best response to a given set of allocation barriers.

The notions of allocative efficiency and allocation barriers are important because they open the door for unilateral policy to actively influence global transport cost, emissions and welfare without imposing a direct tax on transport emissions. Instead, global emissions can solely change in response to a reallocation of ships that is initiated by the imposition or removal of allocation barriers.

Unilateral policies aiming to reduce carbon emissions can therefore either target transport cost directly in form of carbon taxes or do so indirectly by interfering with the allocation process. To understand the effectiveness of each approach, we need a theoretical framework which links trade, transport and emissions which is presented in the next section.

4 Theoretical Framework

In this section, I present a theoretical framework in which trade is routed indirectly though an endogenous transport network creating carbon emissions in the process. In the model, different production technologies across countries create a Ricardian motive for trade and the indirect nature of shipments necessitates the choice of an optimal route between each origin and destination. Both features are achieved by combining a generalized Eaton-Kortum framework of trade with the route selection model of Allen and Arkolakis (2022). The key novelty is to express transport cost as the outcome of an endogenous allocation process in the shipping sector. Operators are tasked with finding an optimal allocation of ships across a fixed set of circular routes and need to serve transport demand at the

lowest cost possible. Cost incentives generated by heterogeneity in the global fleet lead to an allocation in which traffic-intensive network segments experience higher levels of fuel efficiency and lower transport cost. The allocation offers a new microfoundation for scale effects in the shipping industry, endogenously responds to changes in transport demand and thereby provides a unified framework to study interactions between trade, transport and emissions. In section 4.3, I use this framework to study the impact of unilateral policy on transport emissions. An estimation of model parameters and quantification of counterfactual scenarios are deferred to sections 5 and 6.

4.1 Setup

Consumption and Production: The baseline setup closely follows Ganapati et al. (2024) (GWZ) who introduce route selection into a general equilibrium model of trade to study the maritime transport sector. In an Eaton and Kortum (2002) world with J countries and N industries, consumers in each country j choose from continuum of goods ω_n which are homogenous within, but heterogeneous across industries n.[18] Consuming goods grants consumers utility according to function $U_j = U_j(C_j)$ that is continuous and twice differentiable. The function's argument C_j is a matrix capturing the chosen quantity of each good consumed in j.

On the production side, goods in sector n of country i are produced according to production function $q_{in}(\omega) = f_{in}(z_{in}, L_{in}, Q_{in})$ where $f_{in}(.)$ is continuous and twice differentiable, z_{in} represents the production technology, L_{in} is a vector of non-tradable factor inputs and Q_{in} is a matrix of goods from other industries used as inputs in production of good ω_n. The production function and all inputs vary at the country-sector level but are common within industries of the same country. With perfect competition among firms, the marginal cost of goods in a country-sector is

$$c_{in} \equiv c_{in}(z_{in}, W_i, P_i)$$

[18]The data used in this paper allows to define industries at a highly disaggregated level to ensure the assumption holds for all sectors of the economy.

where W_i is a vector of factor prices in country i and P_i a matrix of prices of all goods used as inputs in industry n.

Pricing: Goods can be sourced domestically or purchased from foreign markets. The price of a good produced in country i and sold in country j is

$$p_{ijn}(\omega) = c_{in}\kappa_{ijn}\tau_{ijnr}(\omega)$$

It consists of marginal cost c_{in}, bilateral trade costs such as tariffs which are unrelated to transport κ_{ijn} and transport costs τ_{ijnr} incurred when shipping ω along the lowest cost route r connecting country i and j. In standard Eaton-Kortum fashion, goods are purchased if the combined cost of production, trade and transport undercuts the price of all other goods from the same sector offered in country j. This Ricardian selection provides a motive for trade and ensures that observed prices p_{ijn} represent least-cost combinations in all markets.

Shipping: A key distinction from canonical models of trade is the presence of a transport network. Rather than being forced to transport goods directly, producers can choose an indirect route involving an arbitrary number of intermediary stops K_r before reaching their destination. I use the term *network leg* or *network segment* interchangeably to describe a direct link connecting two stops. Following Allen and Arkolakis (2022), transporting goods over a network segment belonging to route r creates leg cost

$$t_{k_{r-1},k_r} = f(\Xi, \varepsilon_{k_{r-1},k_r}) \tag{1}$$

which is a function of endogenous network traffic Ξ and exogenous cost components $\varepsilon_{k_{r-1},k_r}$ such as leg distance. Leg costs t_{k_{r-1},k_r} are expressed ad valorem. Chaining up all segments used along a route delivers a deterministic route cost

$$\tilde{\tau}_{ijr} = \prod_{k=1}^{K_r} t_{k_{r-1},k_r}(\Xi, \varepsilon_{k_{r-1},k_r})$$

Producers seeking to minimize transport cost can therefore compare different direct and indirect routes $\tilde{\tau}_{ijr}$ connecting origin i and destination j.[19] To allow variation in route choices within origin-destination pairs, route costs are subject to an idiosyncratic shock drawn from a Fréchet distribution with cumulative distribution function (cdf)

$$F_{ijn}(\epsilon) \equiv \Pr\{\epsilon_{ijnr}(\omega) \leq \epsilon\} = \exp\left\{-\epsilon^{-\theta}\right\} \tag{2}$$

where shape parameter θ governs the strength of dispersion in route choices between origin i and destination j. The stochastic transport cost[20] of a route is

$$\tau_{ijnr} = \frac{\tilde{\tau}_{ijr}}{\epsilon_{ijnr}(\omega)} \tag{3}$$

The presence of idiosyncratic route shocks rationalizes deviations from deterministic least-cost routes $\tilde{\tau}_{ijr}$ connecting i and j and matches the observed dispersion in route choices in the data.

The gravity of traffic: Ricardian selection implies that the observed price of good ω shipped along route r between countries i and j must represent the lowest price among all competitors supplying the same good along different routes r' or from different origins i'. Assuming Fréchet distributed route shocks, the joint probability that $p_{ijnr}(\omega)$ is the lowest price among all other offers can be expressed as

$$\pi_{ijnr}(\omega) \equiv \Pr\left\{p_{ijnr} \leq \min_{i' \in I \setminus i, r' \in R_{ij} \setminus r} p_{i'jnr'}\right\} = \frac{\left[c_{in}\kappa_{ijn}\tilde{\tau}_{ijr}\right]^{-\theta}}{\sum_{i' \in I}\left[\left(c_{i'n}\kappa_{i'jn}\right)^{-\theta} \cdot \sum_{r' \in R_{i'j}} \tilde{\tau}_{i'jr'}^{-\theta}\right]} \tag{4}$$

as shown in appendix C.

As each route is chosen with non-zero probability, it is useful to derive an expression which aggregates the transport cost of all routes taken by industry n supplying country j from country i. To do so, I first introduce adjacency matrix $A_n = \left[t_{kln}^{-\theta}(\Xi, \varepsilon_{kl})\right]$ which captures the direct cost of traversing a single segment of the transport network, scaled by Fréchet parameter θ. The total transport cost of all routes of any length between i

[19] Note that $\tilde{\tau}_{ijr}$ is common across all goods choosing route r.

[20] Note that higher draws of $\epsilon_{ijnr}(\omega)$ lead to a reduction in route costs τ_{ijnr}

and j can then be expressed by matrix B_n which is the Leontief inverse[21] of the weighted adjacency matrix A_n:

$$B_n(\Xi, \varepsilon) = [\tau_{ijn}] \equiv \sum_{r \in R_{ij}} \tilde{\tau}_{ijr}^{-\theta} \equiv \left[(I - A_n(\Xi, \varepsilon))^{-1}\right]^{\circ(-\theta)},$$

where $\circ$ represents an element-wise Hadamard power. Allen and Arkolakis (2022) show that expressions of direct (t_{kl}) and total (τ_{ij}) transport cost collected in matrices $A_n(\Xi, \varepsilon)$ and $B_n(\Xi, \varepsilon)$ can be used to describe the relevance of individual links in the transport network. They define link intensity π_{ijn}^{kl} as the expected number of times a segment (k, l) is used for bilateral trade between countries (i, j). Summing across all routes connecting country pairs (i, j) and using sourcing probability π_{ijnr} from equation 4, link intensity can be expressed[22] entirely as a function of elements of cost matrices A_n and B_n:

$$\pi_{ijn}^{kl} = \sum_{r \in R_{ij}} \frac{\pi_{ijnr}}{\sum_{r' \in R_{ij}} \pi_{ijnr'}} n_r^{kl} = \left(\frac{\tau_{ijn}}{\tau_{ikn} t_{kln} \tau_{ljn}} \right)^{\theta} \tag{5}$$

Intuitively, the second representation of link intensity in equation 5 compares the cost of the optimal routes chosen by producers, τ_{ijn}, to the cost of alternative routes connecting (i, j) via segment (k, l). If all least-cost routes naturally contain link (k, l), both costs will be identical and the link intensity equal to one.[23] If routing trade via segment (k, l) however raises route costs beyond τ_{ijn}, link intensity π_{ijn}^{kl} will decrease, reflecting a lower relevance of the link for trade between (i, j).[24]

Multiplying link intensity from equation 5 with bilateral trade X_{ij} and summing across all country pairs (i, j) yields the equilibrium volume of traffic flowing along each network segment (k, l):

$$\Xi_{kln} = \sum_i \sum_j X_{ijn} \left(\frac{\tau_{ijn}}{\tau_{ikn} t_{kln} \tau_{ljn}} \right)^{\theta} \tag{6}$$

[21] Representing matrix B_n as a Leontief inverse of adjacency matrix A_n requires the spectral radius of A_n to be less than one. This is the case if $\sum_j t_{ijn}^{-\theta} < 1$ for all i.

[22] n_r^{kl} describes the number of times route r passes through link (k, l). The derivation steps for equation 5 are shown in appendix A.2. of Allen and Arkolakis (2022).

[23] Note that producers are assumed to choose routes optimally, meaning τ_{ijn} capture the lowest possible transport cost between countries i and j. Link intensity is therefore restricted to the unit interval.

[24] Forcing US-China trade (i, j) to be routed via New Zealand and Australia (k, l) would add considerable cost compared to a direct service connecting both countries and result in a low link intensity.

In this gravity equation, segment traffic Ξ_{kl} is directly related to origin-destination trade X_{ij} and elements of cost matrices A_n and B_n which together determine link intensity. As link intensity is invariant to changes at the origin-destination level, other factors such as tariffs or multilateral resistance have no impact on segment traffic after conditioning on bilateral trade.[25]

Equation 6 provides a direct link between trade, traffic and the transport network. As such, it represents an ideal starting point to study the propagation of shocks caused by unilateral changes in policy. To assess the environmental impact of such policies, the next section introduces carbon emissions by modelling the optimal behaviour of ship operators in the transport sector.

4.2 Transport sector

Scale effects: Throughout the previous section, transport cost matrices A and B are assumed to be endogenous and to vary with the current level of network traffic Ξ. While the model imposes no restrictions on the relationship between transport volume and cost, descriptive evidence presented in section 3 and empirical estimates in the literature[26] all suggest the presence of scale economies in the shipping industry which manifest in a negative relationship between cost and volume ($dt_h/d\Xi_h < 0$).

In this paper, scale effects emerge as the result of an optimal allocation of transport technology in the shipping sector. The key assumption maintained throughout the paper is that transport cost t_h is a function of the average fuel efficiency $\bar{F}_h$ of all ships allocated to network segment h[27]

$$t_h = f(\bar{F}_h) \quad \text{where} \quad \frac{\partial t_h}{\partial \bar{F}_h} > 0 \tag{7}$$

[25] A reduction in tariffs for example leads to proportional changes in cost across all routes. The implicit comparison between optimal and alternative routes shown in equation 5 therefore remains unchanged.

[26] Ganapati et al. (2024) find that a 1% increase in traffic reduces leg costs by -0.06%. At the origin-destination level, Asturias (2020) estimates a trade-cost elasticity of -0.23.

[27] To facilitate the discussion, network segments from now on are indexed using edge notation h rather than node notation (k, l).

where function f is continuous, twice-differentiable and increasing in fuel efficiency.[28] The intuition behind this assumption is simple. Segments receiving a more efficient allocation of ships experience lower levels of unit fuel consumption $\bar{F}_h$. Unit fuel savings are then passed on to consumers in the form of lower transport cost t_h. The allocation of ships across the network therefore directly determines transport cost on every segment.

As the allocation is an endogenous outcome that responds to changes in transport volume Ξ_h, we can explain scale effects in maritime shipping as the product of two distinct elasticities: a positive elasticity governing the pass-through of fuel savings to transport cost, and a negative elasticity describing the allocation behaviour of vessel operators

$$\underbrace{\frac{dt_h}{d\Xi_h}}_{\text{total scale effect}} = \underbrace{\frac{\partial t_h}{\partial \bar{F}_h}}_{\text{pass-through}} \cdot \underbrace{\frac{\partial \bar{F}_h}{\partial \Xi_h}}_{\substack{\text{fuel elasticity} \\ \text{(stylized fact 3)}}} \tag{8}$$

As shown in stylized fact 3, unit fuel consumption decreases with segment traffic rendering the product of both elasticities negative. The origin of scale effects in maritime shipping can therefore be traced back to the same culprits responsible for the allocation of technology: heterogeneity in the global fleet (stylized fact 1) and a positive correlation between load factors and traffic (stylized fact 2).

Linking transport cost to the allocation of technology provides a new microfoundation for scale effects and distinguishes this paper from the existing endogenous transport network literature.[29] In this paper, transport cost, scale and emissions all endogenously change with the allocation of technology which is described next.

Allocation setup: There are two important features of the container shipping industry that allow me to model the vessel allocation in a tractable manner: an inelastic supply of transport services that is plagued by overcapacity and a stable set of container routes that are directly observed in the data.

[28] This assumption is formally tested in appendix D.5. Results using external price data suggest a pass-through rate 1.115.

[29] Ganapati et al. (2024) also estimate leg-level scale effects the maritime transport network but resort to an estimation approach that remains agnostic about the underlying mechanism. This forces them to use a fixed scale parameter in each counterfactual whereas scale effects in this paper change endogenously with the allocation of technology.

First, the supply of transport services is almost completely inelastic in the short to medium run because lengthy construction times and high investment costs prevent shippers from a strategic expansion or withdrawal of capacity in response to changes in transport demand.[30] At the same time, persistent overcapacity caused by long-run investment cycles[31] forces shippers to price close to marginal cost.[32] In this paper, I therefore assume that transport capacity is quasi-fixed[33] and that any variation in transport cost is related to changes in transport demand[34] via the optimal allocation of transport technology as shown in equation 8.

Second, the majority of containerized traffic occurs on a stable set of circular routes L as shown in appendix A.2. Ships serve demand by passing through a fixed sequence of route segments (L, h) before returning to their original point of departure. This traffic pattern distinguishes container shipping from other forms of maritime transport[35] and greatly facilitates the allocation procedure as demonstrated below. I treat route formation as exogenous which means that unilateral policy can affect traffic across routes but leaves

[30] Evidence supporting an inelastic supply of transport services is presented in appendix B.3. Figure 13a shows that the average construction time of a modern container ship is three years. Operators therefore cannot respond to a sudden surge in transport demand by expanding their fleet. At the same time, 96% of all ships are active at any given time as shown in Figure 13b. Paired with an average vessel service life of 20 years, this indicates that the large initial investment cost forces operators to continuously deploy all available vessels over an extended period of time to achieve any return on investment.

[31] Leibovici and Dunn (2023) show that investment in new ships concentrates around rare demand spikes such as the global financial crisis or the Covid-19 pandemic. Due to capacity constraints in the shipbuilding industry, global transport supply increases steadily in the decade following the shock, causing tremendous overcapacity after the initial demand shock subsides.

[32] Average operating margins of shipping companies in the sample period (2014-2018) are close to zero (Notteboom et al., 2022). South Korean shipping company Hanjin even declared bankruptcy in 2017.

[33] Specifically, I assume that the *level* of transport capacity cannot increase and that the *share* of available clean and dirty capacity remains unchanged. Conversely, I do allow shippers to withdraw capacity to avoid a capacity surplus and consider alternative fleet compositions in appendix E.3.

[34] This reduces the objective of vessel operators to a simple cost-minimization process in which they choose an allocation that serves a given level of transport demand at the lowest possible cost. This assumption is appropriate for a setting in which global transport demand never exceeds total transport capacity. The counterfactuals studied in this paper all lead to a reduction in transport demand which ensures that the transport sector never operates at capacity. Alternative settings in which demand exceeds capacity such as in the aftermath of the global financial crisis in 2009 or the Covid-19 pandemic in 2021 would require a different set of assumptions.

[35] Bulk carriers studied by Brancaccio et al. (2020) for example operate similar to taxis in that they travel wherever demand (or the anticipation thereof) takes them. Container ships on the other hand are more comparable to buses which follow a predetermined route and (un)load containers along the way.

the set of available routes unchanged.[36]

Optimal allocation: Containers are transported along a predetermined set of routes $\mathcal{L}$. Each individual route $L \in \mathcal{L}$ consists of a fixed sequence of route-segments (L, h) of length $d_{L,h}$. Shippers operating these routes face a given level of transport volume $\Xi_{L,h}$ on each segment and have to allocate a sufficient amount of transport capacity to the route to ensure that all containers are transported.

Transport capacity is provided by two[37] types of ships s which differ in size N_s, per-km fuel consumption $\hat{\delta}_s$ and fuel efficiency $\delta_s = \hat{\delta}_s/N_s$. Following stylized fact 1, I assume that

A1: The unit fuel consumption of clean ships, denoted δ_c, is smaller than the unit fuel consumption of dirty ships δ_d such that $\delta_c < \delta_d$ (= between ship heterogeneity).

A2: The unit fuel consumption of any ship type s decreases with load factor μ such that $\partial\delta_s(\mu)/\partial\mu < 0$ (= within-ship heterogeneity).

A3: Clean ships are larger than dirty ships such that $N_c > N_d$.

The available amount of clean capacity B_c is strictly smaller than total transport demand. Shippers therefore have to decide how to optimally allocate their clean capacity to serve transport demand on all routes at the lowest cost possible.

To solve this complex allocation problem, I exploit the fact that the minimum amount of capacity a circular route needs to receive is uniquely determined by the route-segment with the largest trade volume $\Xi_{L,max}$. A ship of size N_s requires $(\Xi_{L,max}/N_s)$ route rotations to serve demand on all segments which means the entire route consumes $\sum_h \Xi_{L,max} d_{L,h}$ container-km of transport capacity. As the number of route rotations is determined by $\Xi_{L,max}$, ships travelling over the top segment will always be fully loaded ($\mu = 1$). On all other segments, ships traverse a segment at full load for the first $(\Xi_{L,h}/N_s)$ rotations but do not carry any containers on the remaining trips. The fuel expenditure of any

[36] Between 2016 and 2018, 76%-82% of annual transport volume occurs on routes that are also available in an adjacent year. I therefore leave an endogenous route formation for future work.

[37] Note that the framework can accommodate any number of ship types as long as vessel efficiency δ_s increases with vessel size N_s.

route-segment (L, h) can therefore be written as

$$F_{L,h} = \Bigg(\underbrace{\hat{\delta}^{\mu=1}_{L,s} \cdot \frac{\Xi_{L,h}}{N_s}}_{\text{cost of full trips}} + \underbrace{\hat{\delta}^{\mu=0}_{L,s} \cdot \frac{\Xi_{L,max} - \Xi_{L,h}}{N_s}}_{\text{cost of empty trips}} \Bigg) \cdot d_{L,h} \tag{9}$$

where notation $\hat{\delta}^{\mu}_{L,s} = \hat{\delta}_{L,s}(\mu)$ refers to the per-km fuel consumption of vessel s travelling under load factor μ. Summing across route-segments, the total route cost can then be expressed as

$$F_L = \delta^{\mu=1}_{L,s} \cdot \sum_h \Xi_{L,h} d_{L,h} + \delta^{\mu=0}_{L,s} \cdot \left[\sum_h \Xi_{L,max} d_{L,h} - \sum_h \Xi_{L,h} d_{L,h} \right]. \tag{10}$$

To determine how to allocate the limited amount of clean capacity, shippers compare the route cost incurred under each ship type. The savings associated with replacing all dirty ships on a route with clean ships are

$$R_L = \left\{ \left(\delta^{\mu=1}_{L,c} - \delta^{\mu=1}_{L,d} \right) + \left(\delta^{\mu=0}_{L,c} - \delta^{\mu=0}_{L,d} \right) \cdot \left[\frac{\sum_h \Xi_{L,max} d_{L,h}}{\sum_h \Xi_{L,h} d_{L,h}} - 1 \right] \right\} \cdot \sum_h \Xi_{L,h} d_{L,h} \tag{11}$$

and consist of two components: total route demand $\sum_h \Xi_{L,h} d_{L,h}$ and savings per container-km represented by the expression in curly brackets.

While savings generally increase with route demand, they can differ across routes even if demand on all routes is equal. The reason for this interesting finding is related to the level of demand dispersion within each route. Represented by the term in square brackets, demand dispersion is high if the majority of route demand is concentrated on the top segment (L, max). High dispersion increases the number of empty trips and raises the overall cost of operating the route. Shippers comparing two routes with the same demand but different levels of dispersion will therefore always prefer to allocate clean capacity to the route where dispersion is higher.

Equation 11 allows me to derive a unique allocation rule for the entire transport network that minimizes transport cost. For any given vector of transport demand $\Xi_{L,h}$, set of routes $\mathcal{L}$ and clean budget B_c, routes can be ranked according to their savings potential R_L. Clean capacity is then allocated sequentially according to the ranking until the budget

is depleted. The optimal allocation rule for clean capacity then becomes

$$\text{rank}(R_L) = r \quad \Rightarrow \quad \delta^*_{r,s} = \begin{cases} \delta_{r,c} & \text{if } \sum_{1}^{r}\left(\sum_{h} \Xi_{L,max} d_{L,h}\right) \leq B_c \\ \delta_{r,d} & \text{if } \sum_{1}^{r}\left(\sum_{h} \Xi_{L,max} d_{L,h}\right) > B_c \end{cases} \tag{12}$$

where $\delta^*_{r,s}$ denotes the technology type s allocated to route L with rank r.

Surplus capacity: After all clean ships have been allocated according to equation 12, demand on remaining routes is served by dirty vessels. Using L' to denote the set of routes that do not receive any clean capacity, the total dirty capacity required to serve demand is $B'_d = \sum_{L'} \sum_h (\Xi_{L',max} d_{L',h})$. Given the persistent overcapacity in the container shipping industry, I assume that there is always a sufficient amount of dirty capacity B_d available to serve demand on routes L'. While this ensures that demand is cleared on all routes, it raises the question of how shippers deal with a potential capacity surplus in case $B_d - B'_d > 0$.

Running a capacity surplus increases total transport cost and emissions because any route that accommodates additional vessels will experience lower load factors which decreases fuel efficiency $\delta_d(\mu)$. Under an inelastic supply of transport services, all surplus capacity has to be allocated to the network. To relax this assumption, I allow operators to idle a fixed share η of their dirty fleet to reduce the surplus capacity to $(1-\eta) \cdot B_d \geq B'_d$ and present results for different levels of η in section 6.[38]

If idle rates are not sufficient to eliminate surplus capacity, we need a second allocation rule to determine how the surplus should be optimally allocated across routes. To do so, I compute the marginal increase in route cost from an additional unit of surplus capacity. The derivation is presented in appendix C.2. Assuming a common maximum load factor $\mu_{L',max} < 1$ across all routes L', the cost increase per unit of surplus capacity is

$$C_{L'} = (\delta^{\mu=\mu_{L',max}}_{L',d} + \beta \mu_{L',max}) \cdot \bar{\mu}_{L'} \quad + \quad \delta^{\mu=0}_{L',d} \cdot (1 - \bar{\mu}_{L'}) \tag{13}$$

[38] Historically, operators have responded to negative demand shocks by idling a share of their fleet or increasing scrapping rates to limit cost increases from overcapacity. The largest recorded fleet reduction from both measures over the past two decades is 12% (Monios, 2023). An increase in surplus capacity beyond this rate is therefore unlikely to be offset by a deliberate reduction in capacity.

where $\bar{\mu}_{L'} = \frac{\sum_h \Xi_{L',h} d_{L',h}}{\sum_h \Xi_{L',max} d_{L',h}}$ represents the average load factor on route L' and $\beta = -(\mathrm{d}\delta_{L',d}/\mathrm{d}\mu)$ is the within-ship load factor elasticity.

Equation 13 indicates that the cost increase from surplus capacity is inversely proportional[39] to average load factor $\bar{\mu}_{L'}$. To minimize total cost, shippers should therefore always allocate their remaining vessels to routes where average load factors are high.

Routes receiving surplus capacity experience a decrease in load factors on all segments which increases total capacity supplied to the route by factor $(1/\mu_{L',max})$ as shown in appendix C.2. Ranking routes L' according to equation 13, the allocation rule for surplus capacity can thus be summarized as[40]

$$\text{rank}(C_{L'}) = k \quad \Rightarrow \quad \mu^*_{k,max} = \begin{cases} < 1 & \text{if } \sum_1^k \left(\frac{1}{\mu_{k,max}} \sum_h \Xi_{L',max} d_{L',h} \right) \leq (1-\eta) B_d - B'_d \\ = 1 & \text{if } \sum_1^k \left(\frac{1}{\mu_{k,max}} \sum_h \Xi_{L',max} d_{L',h} \right) > (1-\eta) B_d - B'_d \end{cases} \tag{14}$$

where $\mu^*_{k,max}$ denotes the maximum load factor on route L' obtaining rank k under the optimal allocation of surplus capacity.

Together, the deterministic allocation rules specified in equations 12 and 14 characterize the optimal allocation of clean, dirty and surplus capacity for any combination of ($\Xi_{L,h}$, $\mathcal{L}$, B_c, η). Reallocations caused by a change in transport demand $\Xi_{L,h}$, barriers imposed on a subset of routes $\mathcal{L}$, an expansion of clean budget B_c or a more elastic supply of transport services η can all be nested in this framework, making it a useful tool to study policy counterfactuals.

Allocation outcome: Collecting the allocation outcomes of equations 12 and 14 in tuple $(\delta^*_{L,s}, \mu^*_{L,max})$ for every route L in the network, we can describe the fuel cost on route-segments under the optimal allocation as

$$F^*_{L,h} = \frac{1}{\mu^*_{L,max}} \left(\delta^{\mu=\mu^*_{L,max}}_{L,s} \cdot \Xi_{L,h} d_{L,h} \; + \; \delta^{\mu=0}_{L,s} \cdot (\Xi_{L,max} - \Xi_{L,h}) d_{L,h} \right). \tag{15}$$

[39] This statement requires that $(\delta^{\mu=\mu_{L',max}}_{L',d} + \beta \mu_{L',max}) < \delta^{\mu=0}_{L',d}$ which is always the case for estimates of fuel efficiency δ_c and δ_d, as well as within-ship load factor elasticity β presented in section 5.

[40] In words, equation 14 means that surplus capacity should be allocated to all routes L' ranked by $C_{L'}$ until the total capacity consumed by these routes is equal to total capacity surplus $(1-\eta)B_d - B'_d$.

Dividing by transport demand $\Xi_{L,h}d_{L,h}$ and summing across routes L, we get the average fuel efficiency of network segment h as

$$\bar{F}_h^* = \sum_L \frac{w_{L,h}}{\mu^*_{L,max}} \left(\delta_{L,s}^{\mu=\mu^*_{L,max}} + \delta_{L,s}^{\mu=0} \left[\frac{\Xi_{L,max}}{\Xi_{L,h}} - 1 \right] \right). \tag{16}$$

where weights $w_{L,h} = \Xi_{L,h} / \sum_L \Xi_{L,h}$ capture the share of transport volume $\Xi_{L,h}$ in total segment volume. Equation 16 summarizes the model-implied allocation of technology in the transport network which is fully consistent[41] with the observed allocation presented in Figure 6a.

Global transport emissions from containerized trade can then be calculated as

$$E = \sum_h \left[\bar{F}_h^* \cdot \Xi_h d_h \cdot \xi \right] \tag{17}$$

where ξ is an exogenous emission factor converting fuel to CO2. Equation 17 nicely illustrates that transport emissions are the direct result of the level of transport demand in the network $\Xi_h d_h$ and an optimal allocation of technology which determines the fuel efficiency $\bar{F}_h^*$ on each segment.

4.3 Unilateral policy

Before turning to the data, I use the theoretical framework presented above to illustrate the two main channels[42] through which unilateral policy can shape global transport emissions: unilateral changes in transport cost and unilateral interference in the allocation process. This section showcases the key mechanism underlying each channel and highlights the trade-offs and distortions arising in the process.

Unilateral change in transport cost: Carbon taxes, canal fees and port charges all constitute important cost components that directly respond to changes in unilateral pol-

[41] Higher levels of transport volume $\Xi_{L,h}$ draw clean (surplus) capacity towards (away from) the segment while also lowering fuel consumption related to demand dispersion within connected routes ($\Xi_{L,max}/\Xi_{L,h}$).

[42] The mechanisms governing each channel are not unique to policy interventions. In fact, unexpected supply chain disruptions such as the blockage of the Suez Canal in 2021 by the container ship Evergiven or the Houthi attacks in the Red Sea in 2023 operate through the same channels and can therefore be studied within the same framework. As the goal of this paper is to investigate whether unilateral policy can serve as an active lever of sectoral decarbonization, I nonetheless focus on a set of deliberate policy interventions in the discussion below, leaving an analysis of random supply chain disruptions for future work.

icy. Raising cost t_h on a single network segment alters optimal routing decisions in the entire network and thereby affect the level of traffic Ξ_h on every network segment as shown in equation 6. While the exposure to cost shocks differs across network segments, global transport demand and trade will always decrease in a counterfactual world where transport cost has gone up.

This aggregate effect has important implications for the allocation of clean and dirty capacity in the network. Looking at the optimal allocation of clean capacity in equation 12, lower levels of total transport demand ($\sum_L \sum_h \Xi_{L,max} d_{L,max}$) and a fixed amount of clean capacity B_c imply that more routes can be served by clean capacity than before ($r \uparrow$). Conversely, as clean ships take over routes that were previously served by dirty ships, demand for dirty capacity B'_d falls while the amount of available dirty capacity B_d remains unchanged. This leads to an increase in surplus capacity $(1-\eta)B_d - B'_d$ in the network. Surplus capacity causes a decline in average fuel efficiency by reducing load factors on surplus-receiving routes, unless shippers can balance supply and demand (B_d and B'_d) by idling a sufficiently large share of their fleet ($\eta > 0$).[43]

Policymakers aiming to reduce transport emissions via a unilateral change in transport cost therefore need to consider two important factors. First, the supply of transport services needs to be elastic ($\eta > 0$) to reap the full environmental gains of the policy. If shippers are unable to withdraw sufficient capacity in response to falling transport demand, any emissions savings might be offset by a decline in vessel efficiency. Environmental outcomes are therefore closely linked to overcapacity in the transport industry. Second, emission reductions achieved by an increase in transport cost always come at the expense of lower levels of trade. This creates a trade-off between environmental and non-environmental welfare components. As neither side dominates a priori, we need to bring the model to the data to understand whether a cost-based approach to sectoral decarbonization is welfare-enhancing.

[43] If transport supply is fully inelastic ($\eta = 0$), the cost increase from surplus capacity will always be larger than the savings achieved from an expansion of clean capacity across a larger number of routes. In essence, this is because a unit of surplus capacity causes load factors on surplus-receiving routes to decline while load factors on clean routes remain constant.

Unilateral change in allocation barriers: Environmental regulation, infrastructure and political sanctions represent a second set of traffic determinants which respond to unilateral policy. Unlike cost-based interventions, this type of policy affects traffic by linking network access to vessel characteristics which creates a barrier for the allocation of vessels across the network.[44]

In the model, allocation barriers take the form of spatial constraints for the allocation of transport capacity in equations 12 and 14. While some barriers are consistently imposed throughout the paper, others are added by unilateral policy.[45] Imposing a new allocation barrier can trigger a reallocation of clean and surplus capacity away from the optimal set of routes L^*. For shippers, this constitutes an allocation distortion that prevents them from fully acting upon the cost incentives described in stylized facts 1 and 2. A distorted vessel allocation raises the fuel expenditure of shippers and results in higher transport cost for consumers.

What distinguishes allocation barriers from direct cost interventions, is that they only affect transport cost if the added constraint is binding. Figure 4 illustrates this point in a setting where a barrier is introduced alongside a carbon tax. The barrier takes the form of an efficiency threshold $\bar{F}_{policy}$ which on a subset of network segments h_p prohibits shippers from choosing an allocation that results in $\bar{F}_{h_p} < \bar{F}_{policy}$. The baseline allocation $\bar{F}_h$ is assumed to be policy compliant because traffic on all affected segments h_p exceeds the compliance threshold X_p.

The introduction of carbon taxes raises costs in the entire network, shifting fuel efficiency curve upwards to $\bar{F}'_h$ and increases the compliance threshold to X'_p. As long as traffic on segments h_p remains above the compliance threshold, allocation barriers are non-binding

[44] An example of environmental regulation are emission control areas which prohibit the use of sulphur-intensive fuel oils near coastlines (see Hansen-Lewis and Marcus (2022)). Infrastructure characteristics such as the depth of a port naturally limit where ships can berth. Similarly, shippers have to adhere to any political sanctions preventing them from serving certain destinations. In the wake of Russia's invasion of Ukraine, the EU for example banned the entire Russian merchant fleet from accessing European harbors.

[45] When calibrating the model in section 5, I ensure that clean ships of size N_c cannot be allocated to ports that cannot accommodate ships of that size according to the data. To impose this constraint, the allocation of clean capacity in equation 12 is performed on a subset of routes whose smallest port size $min(N_L) \geq N_c$.

Figure 4: Joint impact of carbon taxes and allocation barriers

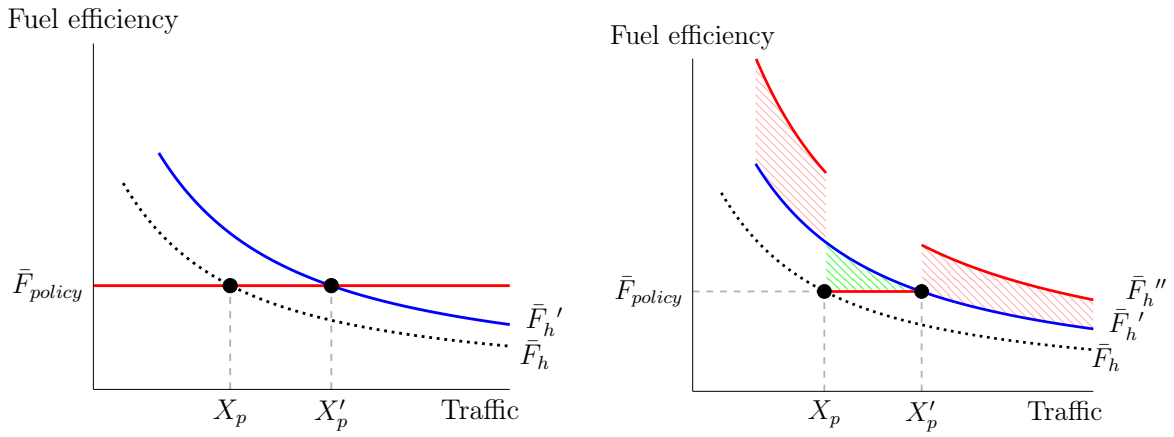

(a) non-binding constraint (b) binding constraint

Note: Figure 4 showcases the impact of carbon taxes and allocation barriers in the theoretical model. It is not based on the observed allocation of vessels and only meant to illustrate the underlying effects.

and have no distinct impact on transport emissions as shown in panel 4a. If traffic on affected segments however falls below X_p', allocation $\bar{F}_h{}'$ is no longer policy-compliant. To re-establish compliance, clean capacity has to be reallocated from other parts of the network towards segments h_p, causing another shift in fuel efficiency from $\bar{F}_h{}'$ to $\bar{F}_h{}''$ as shown in panel 4b. This allocation distortion creates local emission savings on segments h_p at the expense of higher emissions everywhere else and increases transport emissions on aggregate.

Unilateral allocation barriers therefore have no distinct impact on transport emissions if the optimal response of shippers is 'naturally compliant' with the spatial constraint imposed by the policy. This theoretical insight can help us to interpret counterfactual scenarios in section 6 in which both types of unilateral interventions, cost-based and allocation-based, are introduced at the same time.

5 Estimation

I now take the model to the data to estimate key allocation parameters, recover transport costs and estimate the scale elasticity in maritime shipping. The calibrated model can replicate the observed allocation of vessels, matches trade and traffic flows in the network and delivers elasticity estimates that underline the importance of scale economies in the maritime shipping industry.

5.1 Matching the optimal allocation of transport capacity

A key advantage of the theoretical framework in section 4.2 is that it only requires a handful of parameters to capture the entire allocation of vessels across the network. For a given level of network traffic Ξ_h, all we need to know is the heterogeneity in fuel efficiency within and between ships captured by load factor elasticity β and unit fuel consumption of each ship type δ_s.

Within-ship heterogeneity β: Load factor elasticity β is defined as the response in fuel efficiency $\bar{F}_s$ of vessel s to a unit increase in load factor μ. In appendix D.1, I estimate β using hourly fuel consumption data of every container ship in 2018. The load factor elasticity is identified from observed variation in unit fuel consumption within vessels over time as shown in estimation equation 30.[46] Results in Table 6 indicate that raising load factors by 10pp reduces the unit fuel consumption on average between 5.1% and 7.1%. These estimates are robust across a range of ship sizes, load factors, regional restrictions, environmental regulations and daily weather changes as shown in Table 7.

The results indicate that vessel efficiency increases significantly when operating at higher load factors. I incorporate these findings into my model by setting β conservatively at -0.005 which allows the unit fuel consumption $\delta_s(\mu)$ of any vessel[47] to vary by up to 50% such that $\delta_s^{\mu=1} = 0.5 \cdot \delta_s^{\mu=0}$.

Between-ship heterogeneity δ: While parameter β determines the efficiency range *within* vessels, differences in efficiency *between* clean and dirty vessels are captured by parameters δ_c and δ_d. To bring the model as close possible to the data, I create a moment m_1 which captures the difference in total fuel consumption between the matrix of model-implied consumption F^{model} and the matrix of observed fuel consumption F^{data} in the

[46] Estimation equation 30 is $\log(\bar{F}_{s,r,t}) = \beta \cdot \mu_{s,t} + FE_s + FE_r + \varepsilon_{s,r,t}$ where indices (s, r, t) refer to observations at the ship-region-day level. Benchmark estimates therefore account for any unobserved cost factors related to the ship's current state of repair (FE_s) or regional restrictions such as speed limits (FE_r).

[47] Results in Table 6 indicate a slight difference in load factor elasticities across ship types. For simplicity, I assume a common load factor across ship types and choose β conservatively by setting it equal to the lower of the two.

data:

$$m_1(\delta) = \left(F^{\text{model}}(\delta|\boldsymbol{\Xi}^{\text{data}}, \boldsymbol{d}^{\text{data}}, \mu_{max}, N)\right) - \left(F^{\text{data}}\right) \tag{18}$$

where fuel expenditure in the model is a function of vector $\delta = (\delta_c, \delta_d)$, observed traffic and distance matrices $\boldsymbol{\Xi}^{\text{data}}$ and $\boldsymbol{d}^{\text{data}}$, surplus parameter μ_{max}[48] and ship size $N = (N_c, N_d)$. Optimal values of δ_c and δ_d can then be obtained by minimizing moment m_1 using a GMM estimation procedure.[49]

In the benchmark setting, I assume $\mu_{max} = 0.75$, $N_c = 10,000$ and $N_d = 1,000$ which ensures that shippers continue to face an allocation trade-off between dirty and clean capacity across all carbon tax scenarios studied in section 6.[50] The chosen ship sizes further imply that 34% of global transport capacity is provided by clean ships in the baseline.[51] Alternative budget compositions due a reclassification of clean and dirty vessels or investment in clean capacity are explored in appendix E.

Matching outcomes: Matching the model to the observed allocation of capacity proves to be highly successful. Under the optimal efficiency parameters δ_c^* and δ_d^* obtained from equation 18, predicted and observed fuel consumption matrices F^{model} and $F^{\text{model}}(\delta^*)$ show a correlation (in logs) of 0.95 as illustrated in figure 14 in appendix D.2. The model therefore correctly captures the distribution of fuel expenditure across the network.

Beyond fuel consumption F, the calibrated model further matches other important moments that have not been targeted directly by the GMM procedure but are closely linked to allocation behaviour of vessel operators. Key among them are the relationship between load factors and traffic (stylized fact 2) and the relationship between unit fuel consumption and traffic (stylized fact 3) first documented in section 3.

[48] Parameter μ_{max} describes the lowest permissible load factor on the route-segment (L, max). It thereby limits the amount of surplus capacity each route can receive. On routes without surplus capacity $\mu_{max} = 1$ whereas routes with surplus as subject to $\mu_{max} = 0.75$.

[49] To compute model-implied fuel expenditure F^{model}, I aggregate fuel expenditure in equation 15 to the segment level and minimize moment m_1 using an optimal weighting matrix that assigns higher importance to traffic-intensive segments.

[50] For $\mu > 0.75$, all traffic is performed by clean vessels if carbon taxes become sufficiently large. While the model can accommodate this scenario, I defer a discussion of environmental outcomes under a homogenous transport technology to appendix E.3.

[51] To determine the baseline budget shares of clean and dirty capacity, I compute the share of global container traffic (in TEU-km) in 2018 that is performed by vessels with a transport capacity of at least 10,000 TEU.

Figure 5: Matching outcomes

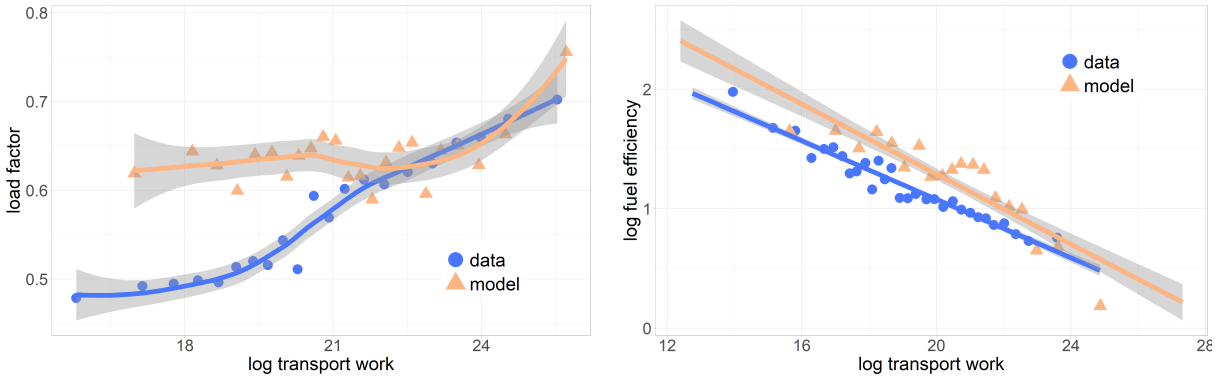

(a) stylized fact 2

(b) stylized fact 3

Note: The figure shows point estimates of binscatter regressions following Cattaneo et al. (2023). Dots refer to estimates based on observed data. Triangles are based on the optimal allocation of transport capacity according to equations 12 and 14 for given levels of traffic.

Figure 5 illustrates how well the calibrated model is able to replicate these stylized facts. Each panel shows results from binscatter regressions of the respective relationship using the observed and model-implied allocation of capacity. In Figure 5a, the model correctly captures the positive correlation between traffic and load factors, albeit more so for segments experiencing medium to high levels of traffic. This is reassuring because it indicates that the model preserves the cost incentive of shippers to allocate their most efficient vessels to traffic-intensive segments. As a result, the allocation of clean capacity across segments closely follows the pattern observed in the data as shown in Figure 5b. Unit fuel consumption in the model decreases with transport work because traffic-intensive segments receive a disproportionate amount of clean capacity.

Despite the underlying complexity of the global transport network, these results indicate that a handful of estimated parameters is sufficient to replicate the most important moments governing the allocation process.

5.2 Retrieving transport cost

While trade and traffic flows are routinely recorded in customs and satellite data as shown in section 2, leg-level transport cost is typically not directly observed in the data. This section explains how transport cost can be recovered from the theoretical framework. I then show how well these estimates can predict observed trade and traffic in the data.

Recovering transport cost: To recover transport cost matrix A, I follow the procedure pioneered by Ganapati et al. (2024) which is explained in detailed in appendix D.3. In short, it exploits model equation 6 which states that the equilibrium flow of leg-level traffic Ξ_{kl} only depends on origin-destination trade flows X_{ij} and leg-level transport cost captured by matrix A. As trade and traffic are observed in the data, we can invert this system to recover a cost matrix that rationalizes Ξ_{kl} and X_{ij}.

To implement this idea empirically, elements of matrix A are predicted using a gravity-style estimation equation $\hat{A} = g(\rho|\mathbf{Y})$ where $\mathbf{Y}$ describes the set of variables used in function g. The goal is to find a parameter vector ρ for which predicted matrix $\hat{A}$, once plugged into model equation 6, returns a predicted traffic matrix $\hat{\Xi}$ that is as close as possible to the observed level of traffic Ξ^{data}. This approach can be summarized in the following moment equation:

$$m_2(\rho) = \left(\hat{\Xi}(\hat{A}(\rho)|\mathbf{X}, \mathbf{Y}, \theta)\right) - \left(\Xi^{\text{data}}\right) \tag{19}$$

where predicted traffic $\hat{\Xi}$ is a function of predicted cost matrix $\hat{A}$, $\mathbf{X}$ represents the observed matrix of trade flows and θ is the trade elasticity.

Predicted transport cost: GMM estimates of parameter vector ρ are reported in Table 19 in appendix D.3. To assess how well recovered transport cost fits the data, I plug matrix $\hat{A}(\rho)$ into equation 6, compute predicted levels of traffic $\hat{\Xi}(\rho)$ and trade $\hat{\mathbf{X}}(\rho)$ and compare them to their observed counterparts Ξ^{data} and $\mathbf{X}^{\text{data}}$. Figure 6 plots the correlation between predicted and observed container traffic and trade in log-shares. Starting with transport volume as the targeted moment in Figure 14a, both measures show a correlation of 0.92. Figure 14b instead studies the correlation between both measures for the untargeted moment, containerized trade, which remains high at 0.68.

The strong correlation across both measures indicates that the recovered cost matrix can indeed rationalize observed levels of traffic and trade. This not only lends credibility to the recovered cost estimates themselves, but also underlines the ability of the model to reliably capture observed behaviour of economic agents.

Figure 6: Evaluating predicted transport cost

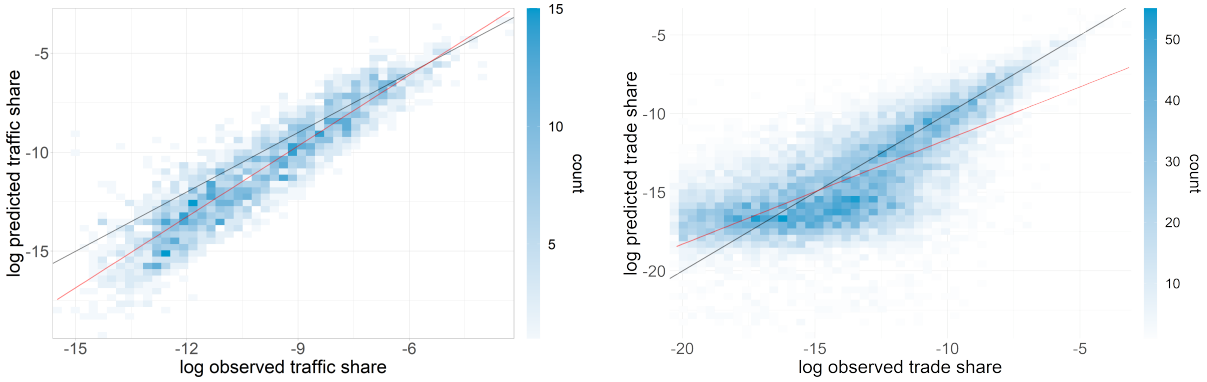

(a) transport volume

(b) containerized trade

Note: The figure plots predicted values of transport volume and containerized trade (y-axis) against their empirical counterparts (x-axis) in the data. Predicted values are obtained by plugging transport cost $\hat{A}$ recovered from equation 19 into equation 6. Each measure is represented as log-shares. The identity line (black) represents a perfect correlation across both measures while the red line captures the actual correlation between the two.

5.3 Estimating aggregate scale effects

Equipped with transport cost (section 5.2) and fuel efficiency (section 5.1), I can now use the theoretical framework to estimate aggregate scale effects in the maritime transport network.

Aggregate scale effects λ_1 describe a decrease in transport cost for higher levels of traffic. In this paper, they arise as the product of two elasticities: a fuel elasticity α_1 capturing the relationship between unit fuel consumption and traffic (= stylized fact 3), and a pass-through elasticity σ_1 governing the relationship between fuel efficiency and transport cost (equation 8).

To compare model-based estimates to existing scale effect estimates in the literature, I proceed in three steps: First, I replicate a model-agnostic approach of GWZ using data from 2018. In this approach, scale effect λ_1 is estimated directly using transport cost recovered in section 5.2 and observed levels of traffic, and therefore does not rely on the allocation mechanism described in section 4.2. Next, I use the framework to derive a theory-consistent equivalent of GWZ and take it to the data. This approach retains the same identification strategy but decomposes the aggregate scale effect into fuel and pass-through elasticities in line with the model. Finally, I present scale effects which are

purely based on the calibrated model. A detailed description of each step is presented in appendix D.3.

Table 1 summarizes the results of each approach at the segment and origin-destination level. The table hereby allows comparisons across two groups. The first group consists of columns 1 and 2 which both apply the estimation procedure of Ganapati et al. (2024) but rely on data from different years. Estimating scale effects from a larger and more recent sample[52] of traffic flows (column 2) leads to a slightly larger scale elasticity than reported in the original paper (column 1). According to this estimate, a 10 percent increase in transport volume on segment h decreases transport cost by roughly 1 percent. The second group consists of columns 2, 3 and 4 which all rely on the same sample but follow different estimation approaches. Comparing a model-agnostic approach (column 2) to a theory-consistent (columns 3) or purely model-based (column 4) estimation procedure does not alter the overall results in a meaningful way. Estimates obtained from a decomposition of scale effects into fuel and pass-through elasticities are therefore consistent with prior results of the transport literature.[53]

When conducting policy counterfactuals, scale effects endogenously change with the allocation of transport capacity to capture changes in vessel efficiency under the new policy environment. This marks an important departure from the literature which typically treats scale effects as an exogenous parameter. Ignoring the endogenous response of shippers to a change in transport demand could introduce severe bias to the estimated environmental gains of the policy. Assuming a constant scale elasticity might for example overstate the negative impact of carbon taxes on bilateral trade, if the reallocation of vessels results in an overall decrease in fuel efficiency. To quantify the importance of this channel, I will directly compare emission estimates under a constant and endogenous scale elasticity when conducting counterfactuals in the next section.

[52] Ganapati et al. (2024) estimate scale effects from a sample containing 6 months of traffic data in 2014 while the replication in this paper is based on all containerized traffic in 2018.

[53] Model-based estimates in column 4 are almost identical to causal estimates obtained from the model-agnostic procedure developed by GWZ. They also remain very close to origin-destination level estimates reported by Asturias (2020) which estimate an elasticity of trade cost to trade volume of 0.26.

Table 1: Scale effect in maritime shipping - summary

Approach:	Ganapati et al. (2024)	replication	observed $\bar{F}_h$	model-generated $\bar{F}_h$
Transport measure:	TEU	TEU	TEU	TEU-km
Scale effect result				
leg-level	-0.06	-0.108	-0.139	-0.106
origin-destination level	-0.17	-0.259	-0.334	-0.254
IV	yes	yes	yes	no

Note: The table summarizes scale elasticity estimates obtained from equations 34-37 following the procedure explained in appendix D.4 and presented in detail in table 9. Scale effects indicate how the transport cost on network segment h responds to a unit increase in transport volume (columns 1-3) and transport work (column 4). Column 1 reports reduced-form estimates from Ganapati et al. (2024). Column 2 replicates their approach using traffic data from 2018. Column 3 instead shows estimates from a theory-consistent estimation equation that is derived in appendix D.4 and applied to observed fuel efficiency $\bar{F}_h$. Column 4 then presents scale effect estimates from the calibrated model. Columns 1-3 represent causal estimates after controlling for endogenous changes in transport demand via a geography-based instrument. All columns assume a trade elasticity $\theta = 4$ and an average route length of 2.4. The latter is based on observed US container journeys described in appendix A.2.

6 Counterfactuals

Counterfactual exercises presented in this section serve two distinct purposes: First, they quantify the environmental gains of the EU's carbon policy and thereby showcase its role in reaching sectoral emission targets. Second, they illustrate that the theoretical framework can generate a breadth of novel insights that help us to evaluate, understand and design new approaches to decarbonize maritime shipping in general.

6.1 Counterfactual methodology

Solving the model in changes: To close the model, I adopt the framework of Caliendo and Parro (2015) and group all products into containerized (ct), non-containerized (nct) and non-tradable (nt) sectors ($n \in [ct, nct, nt]$). The model features global input-output linkages via a roundabout production technology that allows all three industries to serve as intermediate inputs for other sectors in domestic (nt) or foreign (ct and nct) markets. The full general equilibrium model is presented in appendix C.3.

I solve the general equilibrium model in relative changes by studying how key variables of interest respond to changes in transport cost. For this purpose, I adopt the exact-hat algebra notation and describe changes in any variable as $\hat{v} = \frac{v'}{v}$ where v' and v denote counterfactual and baseline values respectively. The primary interest of this paper is to

understand how unilateral policy in the form of carbon taxes or fuel standards affects global emissions and welfare. To this end, I consider policy-induced changes in leg-level transport cost $\hat{t}_{kl} = t'_{kl}/t_{kl}$ and study the endogenous response[54] of trade flows, traffic volumes and carbon emissions.

Additional data: To quantify the full model, I complement traffic and fuel consumption data described in section 2 with bilateral trade data of the BACI database (Gaulier and Zignago, 2010) and retrieve country-level input-output data from the EORA database (Lenzen et al., 2012, 2013). I aggregate trade, consumption and production data to the sector level as required by the model and compute Cobb-Douglas expenditure shares φ and value-added shares γ for each sector. Across all specifications, I follow the literature and set trade elasticity $\theta = 4$ (Simonovska and Waugh, 2014). The final sample contains information for all three sectors across 136 countries in 2018. Further details are presented in appendix A.3.

Counterfactual procedure: Unilateral policy interventions such as carbon taxes or fuel standards raise the cost of transport for a subset of network segments. These leg-level changes in transport cost $\hat{t}_{kl}$ change the expected transport cost $\hat{\tau}_{ij}$ for every origin-destination pair in the network because all of them rely on targeted segments with non-zero probability $\pi_{ij}^{kl} > 0$. While the initial impact of each unilateral policy is local, the consequence of it is a global change in transport cost.

From the perspective of the model, the change in expected transport cost $\hat{\tau}$ is isometric to a change in other bilateral cost factors unrelated to transport such as tariffs. When conducting policy counterfactuals, we can therefore think of each intervention as a highly heterogeneous change in bilateral tariffs whose intensity is determined by the structure of the global transport network.

To initialize leg-level cost changes $\hat{t}_{kl}$, I use benchmark estimates of fuel and pass-through

[54] Changes in leg-level transport cost $\hat{t}_{kl}$ triggers a change in expected transport cost $\hat{\tau}_{ijn}$. The marginal cost of production $\hat{c}_{in}$, price indices $\hat{P}_{in}$, bilateral trade shares $\hat{\pi}_{ijn}$ and wages $\hat{w}_i$ all respond to a change in $\hat{\tau}_{ijn}$ as shown in equations 25-28 in appendix C.3. The resulting system of equations can be solved by finding the vector of wages $\hat{w}_i$ for which counterfactual trade is balanced (up to a deficit shifter) as shown in equation 29 (Caliendo and Parro, 2015).

elasticity from section 5.3 to update local elements of cost matrix A. This unilateral change in transport cost triggers an endogenous response in expected transport cost $\hat{\tau}$, trade flows $\hat{X}$, traffic $\hat{\Xi}$, fuel efficiency $\hat{\bar{F}}$ and scale effects $\hat{\lambda}$ and continues until a new equilibrium is reached. Intuitively, this process can be thought of as a sequence of adjustments in which shippers respond to changes in transport cost by reallocating their vessels to minimize transport cost and consumers in turn respond to updated freight rates by re-optimizing their global sourcing strategy.

To reach a new equilibrium, I follow the algorithm presented in appendix E.1 which ensures convergence[55] to a unique[56] equilibrium as long as scale effects do not exceed the inverse of trade elasticity θ (Kucheryavyy et al., 2023).

6.2 Unilateral carbon policy in action

Unlike other forms of transport, the maritime shipping sector to date is not required to pay for CO2 emissions in any part of the transport network. This situation is about to change. Starting in 2025, the EU has decided to include the shipping sector in its emission trading system (ETS), a cap and trade system in which participants are required to buy allowances for each ton of CO2 emitted within the EU. This unilateral policy applies to 96% of all container ships and covers all direct journeys involving at least one European port, an area accounting for 22% of global container traffic in 2018.[57]

At the same time, the FuelEU Maritime Initiative introduces a new set of fuel standards for vessels operating in EU waters. The policy requires shippers to diversify their fuel mix away from emission-intensive heavy fuel oils (HFOs) and towards other fossil fuels

[55] Every transport cost shock can be represented as change in global tariffs. Once transport costs have been calculated, solving the model therefore follows the same algorithm that was originally developed by Alvarez and Lucas (2007) and is used by Caliendo and Parro (2015) to study the impact of tariff liberalization in context of NAFTA.

[56] This paper focuses on the unique equilibrium that arises from the current equilibrium in 2018. The framework theoretically admits multiple equilibria. Kucheryavyy et al. (2023), however, show that gravity trade models with multiple industries which feature external economies of scale at the industry level remain well-behaved and have a unique equilibrium as long as the product of scale effects and trade elasticity θ remains smaller than 1. Scale effects in this paper are common across industries and on average amount to 0.1 at the leg-level. Based on a sample of US container routes described in section A.2, the average number of legs (at the country level) is 2.4. For a given trade elasticity of $\theta = 4$, this implies that $\lambda \cdot (1/\theta) < 1$.

[57] Ships travelling between two European ports are required to pay for all CO2 emitted on their journey. If journeys involve a non-European port, only half of the emissions require an allowance.

with lower emission factors such as liquefied natural gas (LNG) or its organic counterpart bio-LNG.[58] As only a small subset of vessels which tends to be larger and more efficient is capable of running on LNG, fuel standards can interfere with the vessel allocation by forcing shippers to maintain a minimum share of clean capacity on EU routes.

The introduction of carbon taxes and fuel standards in the EU provide an ideal environment to showcase the advantages of the theoretical framework over existing approaches and investigate the environmental gains of unilateral carbon policy in the shipping sector. To advance in both directions, I create a benchmark scenario in which carbon taxes increase the fuel cost on European legs by 10%[59] without distorting the vessel allocation and shippers can avoid surplus capacity by idling up to 13% of their fleet ($\eta = 0.2$).[60] I then gradually step away from each assumption to highlight important mechanisms of the model.

Aggregate impact of EU carbon taxes: Table 2 summarizes the response to a 10% carbon tax in the benchmark setting. Unilaterally raising the cost of transport on EU legs generates a significant reduction in containerized trade. The magnitude of these estimates underlines the EU's central role in the global transport network and cautions against treating unilateral carbon policy as a regional intervention. While the bulk of trade reductions concentrate in the EU, trade partners, particularly those in the Union's immediate vicinity, experience a similar or even larger decline as shown in appendix E.2. The policy also achieves sizeable reductions in global transport emissions, albeit not in proportion to the reduction in traffic because average unit fuel consumption increases. This presents an important lesson: While taxing maritime transport reduces global traffic

[58] EU regulation 2023/1805 mandates that the energy intensity pooled across the entire fleet of vessels in the EU falls by 6% by 2030, 31% by 2040 and 80% by 2050 relative to a reference value from 2020.

[59] The increase in fuel cost due to an inclusion of maritime shipping in the ETS depends on the price of fuel and emission allowances which vary over time. Initially, shippers only have to account for a share of emissions on affected routes (40% in 2025, 70% in 2026, 100% from 2027). Assuming a fuel price of 600 EUR/tFuel, an allowance price of 65 EUR/tCO_2, an emission factor of 3.114 CO_2/tFuel and using the 2025 emission rate of 40%, fuel costs on EU legs would increase by $(65 \cdot 0.4 \cdot 3.114)/600 = 13.5\%$. The benchmark carbon tax of 10% therefore represents an approximation of the initial impact in 2025.

[60] Parameter η refers to the idle share of dirty transport capacity as shown in equation 14. In the benchmark calibration, dirty capacity accounts for 66% of the global fleet. The implied idle rate of $0.2 * 0.66 = 13.2\%$ is close to the largest recorded fleet reduction of 12% reported in Monios (2023).

Table 2: Benchmark: 10% EU Carbon Tax

change in pp	Global	EU	non-EU
containerized trade	-5.58	-9.18	-1.32
container traffic	-9.48	-12.21	-6.36
fuel efficiency	4.95	12.77	0.14
transport emissions	-7.82	-8.49	-7.20

Note: The table presents counterfactual estimates of a 10% carbon tax on EU segments. The benchmark setting assumes no allocation barriers and an idle rate $\eta = 0.2$.

and thereby transport emissions, it also creates inefficiencies as lower load factors raise the fuel consumption per container-km

Vessel heterogeneity matters for environmental gains: The reduction in fuel efficiency does not eliminate environmental gains of the policy but dampens the fall in transport emissions. Shippers faced with a fall in transport demand on EU segments respond to the average reduction in load factors *within* their fleet and reallocate clean vessels across routes to take advantage of efficiency differences *between* ship types within the new cost environment. Both dimensions of ship heterogeneity are therefore key determinants of the realized environmental gains of the policy.

Existing approaches in the literature have typically failed to establish a link between vessel heterogeneity and environmental outcomes which has severe consequences for the expected emissions savings of unilateral policy. Figure 7 compares the predicted fall in transport emissions in response to a carbon tax of this paper to the state-of-the-art approach in the literature. Ganapati et al. (2024) for example would correctly capture the change in transport demand in the network via estimated scale effects but fail to account for any change in fuel efficiency across segments linked to the reallocation of vessels or lower load factors. The best prediction in this case would be to project the observed level of efficiency on each segment in the baseline onto the counterfactual vector of traffic.

Figure 7 depicts the estimated emission savings of the model (black) to the predicted emission savings under the same vector of traffic but the baseline allocation of vessels (red). Emission savings following the current approach of the literature overestimate the

Figure 7: Homogeneous transport technology overestimates environmental gains

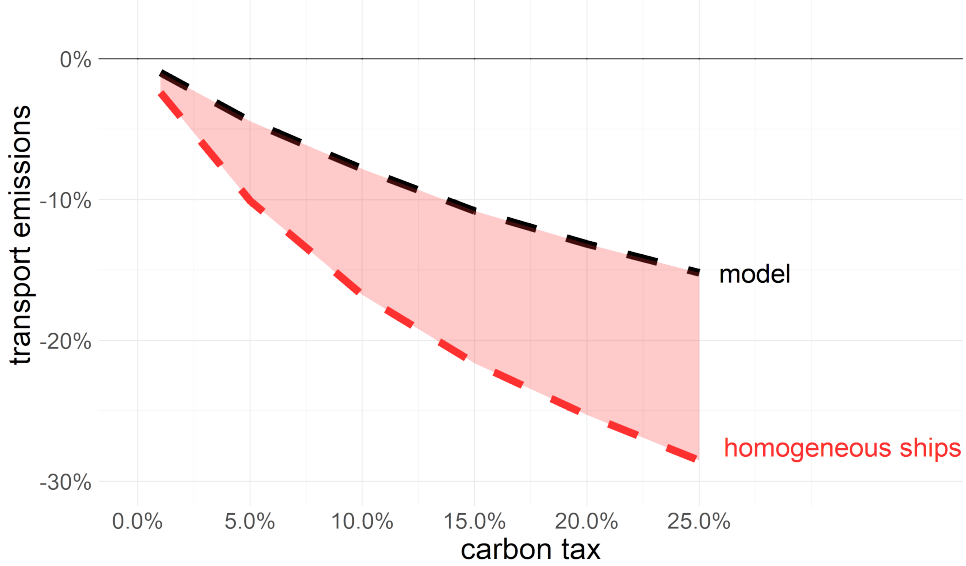

Note: Figure 7 compares counterfactual transport emissions of a model with (black) and without (red) vessel heterogeneity. The former is based on the framework explained in section 4.2 and assumes $\eta = 0.2$ while the latter follows Ganapati et al. (2024). The shaded area indicates the difference of predicted emission savings under both approaches.

environmental gains by a factor of two. Failing to account for vessel heterogeneity by either treating ships as homogeneous (Shapiro, 2016) or fully attributing cost savings to the reduction in traffic (Ganapati et al., 2024) greatly exaggerates the decarbonization potential of unilateral policy.

While accounting for vessel heterogeneity is important to correctly capture environmental gains from unilateral policy, within and between-ship channels are not equally important. To decompose total emission savings into each channel, I re-estimate the full counterfactual model in the baseline setting but prohibit shippers from reallocating vessels across routes[61]. This exercise eliminates any impact of efficiency differences between vessels on transport emissions. Instead, emissions are fully determined by the interaction between transport demand and the load factor elasticity within ships. The exercise reveals that only 14% of baseline emission savings are attributed to the reallocation channel while 86% are linked to changes in fuel efficiency within vessels. Researchers aiming to estimate the

[61]This counterfactual differs from the scenario shown in Figure 7. The graph shows emission savings under the correct vector of traffic but assumes a constant vessel efficiency equal to the baseline allocation in 2018. The counterfactual here instead estimates a new vector of traffic in a setting where shippers cannot reallocate their fleet but fuel efficiency still responds to changes in load factors across segments.

environmental impact of cost shocks in the maritime transport network from observed changes in traffic can therefore approximate counterfactual levels of transport emissions via reduced form regressions as long as they have a precise estimate of load factor elasticity β.

Surplus capacity determines environmental outcomes: The benchmark setting assumes an idle rate $\eta = 0.2$ which implies that shippers can withdraw up to 20% of dirty transport capacity to match transport demand on each route and uphold high levels of efficiency. If shippers instead are unable to adjust capacity to the fall in transport demand ($\eta = 0$) because of rigidities in the supply of transport services, available capacity will exceed transport demand. The resulting surplus then has to be allocated across routes which lowers fuel efficiency on affected segments and raises transport emissions. The environmental response to a unilateral carbon tax therefore crucially depends on the extent to which shippers can respond to a fall in transport demand.

Figure 8 plots the change in transport emissions under an inelastic ($\eta = 0$) and elastic ($\eta = 0.2$) supply of transport services. Allowing for surplus capacity in the transport sector changes the model predictions. Instead of creating emission savings as seen in the benchmark setting, carbon taxes increase global transport emissions above baseline levels. This key finding holds for any carbon tax under consideration. Higher tax rates amplify potential gains and losses but never resolve the ambiguity in environmental outcomes. The same holds for alternative fleet compositions presented in appendix E.3. Increasing the share of clean capacity in the baseline from 35% (benchmark setting) to 45% or 100% never results in any in emission savings as long as the supply of transport capacity remains inelastic.[62]

Carbon taxes are therefore unlikely to generate any environmental gains unless shippers are able to accommodate the rise in surplus capacity. While this insight is generated from a framework that focuses on the short to -medium run, it remains important when

[62] A clean share of 35% in the benchmark is a direct result of the chosen ship size of clean vessels $_c$. The direct relationship between idle rates η and policy outcomes presented here is therefore independent of the chosen definition of clean and dirty vessels

Figure 8: Carbon taxes, transport emissions and surplus capacity

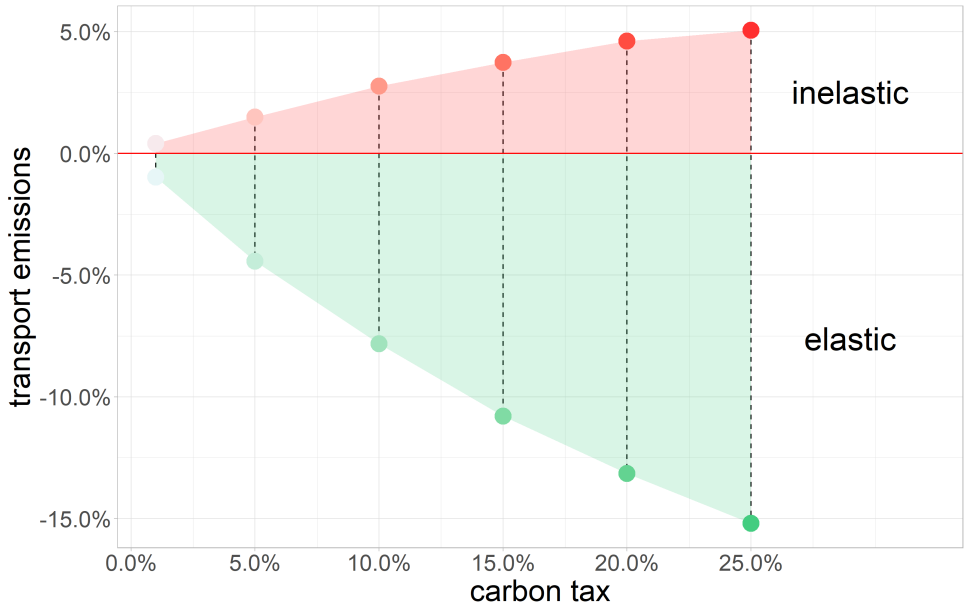

Note: Figure 8 depicts the counterfactual change in global transport emissions for varying levels of carbon taxes and transport supply elasticity η. The graph considers two distinct scenarios. In scenario 1 (red dots), the supply of transport services is fully inelastic ($\eta = 0$). In scenario 2 (green dots), the supply elasticity is equal to its historic maximum observed during the global financial crisis and COVID-19 pandemic ($\eta = 0.2$).

considering that total transport capacity is expected to increase in the long-run. While a formal treatment of a long-run capacity expansion lies outside the scope of this paper, I briefly will return to this point in section 7 and discuss the relationship between surplus capacity and environmental outcomes in light of rising capacity.

Will fuel standards cause an allocation distortion? A final consideration within the benchmark setting is the presence of allocation barriers. Estimating the distortion potential of EU fuel standards is difficult because the data does not allow me to observe which fuel type a ship is using on a given journey in 2018.[63] Rather than trying to quantify the actual fuel standard, I instead simulate a comparable policy scenario in which the EU forces all clean vessels to switch to an alternative fuel type when operating in EU waters. The full details of this counterfactual are presented in appendix E.4. The exercise is meant to resemble the actual policy as closely as possible in that shippers have

[63]Ships equipped with dual-fuel engines could for example switch from HFO to LNG to improve the fuel mix. Without knowing which fuel type they have been using in 2018, it is difficult to compute baseline levels of carbon intensity and assess how much action is needed to meet the new standard.

to ensure sufficient adoption rates of the new fuel type to meet the policy standard but are disincentivized from doing so because the adoption is costly.

The results presented in Table 13 indicate that imposing fuel standards alongside carbon taxes can cause an allocation distortion, but only do so if the policy-induced cost increase from adopting the new fuel type becomes sufficiently large. Fuel standards which only cause a moderate increase in fuel expenditure on EU segments in turn are not severe enough to persuade shippers to reallocate their vessels to other parts of the network. While allocation barriers in general can play a decisive role in shaping environmental outcomes, evidence obtained from simulated EU fuel standards in appendix E.4 suggests that they are unlikely to have a major impact on transport emissions in the short run. For the remainder of the paper, I therefore exclusively focus on the environmental impact of unilateral carbon taxes.

6.3 Social welfare

Consumers respond to endogenous changes in transport cost by re-optimizing their supply chains which alters the locus of production, domestic price levels and wages across countries. To provide a holistic welfare assessment of unilateral policy, we therefore need to look beyond the maritime shipping industry and study the combined effect of transport emissions, production emissions and real wages.

Carbon taxes create environmental spillovers for production: To compute the impact of EU carbon taxes on global production emissions, I combine country-sector specific emission intensity[64] information from the EORA database with the estimated change in trade flows of the structural model. The results are depicted in Figure 16a. A 10% EU carbon tax for containerized shipping is expected to lower global production emissions of traded goods by 1.4pp. On aggregate, I find that that production emissions account for 21% of the total emission savings of the policy. In fact, figure 16b shows that production emissions are the dominant source of emission savings in many countries. A sizeable share of the total environmental gains of the policy are therefore generated outside the

[64]Note that sectoral emission intensity does not vary with the level of production.

transport sector.

EU carbon taxes shift production to emission-intensive countries: While the EU carbon tax causes an overall decline in production emission levels, Figure 16a also indicates that emissions per USD of output have actually increased. This rise in emission intensity implies that carbon taxes on average have shifted the production of traded goods towards country-sectors with dirtier production technology. Figure 9 visualizes this process by showing the change in containerized export shares across countries. Unsurprisingly, European producers suffer the largest decline in exports which benefits competitors in the US, Australia, Canada and China. As these countries on average exceed European emission intensity by 4%, 23%, 30% and 110% respectively, global emission intensity in production increases.

These findings carry two important messages: First, they underline the close relationship between production and transport emissions. Taxing one will inherently affect the other which is why we shouldn't treat them as separate entities when assessing the impact of environmental policies. Second, if carbon policies have to be designed unilaterally, environmental spillovers between both emission types will strongly depend on the chosen target. In the current example, a unilateral intervention in the EU does not realize the full CO2 saving potential because it diverts production towards more emission-intensive

Figure 9: Production of traded goods shifts in response to the carbon tax

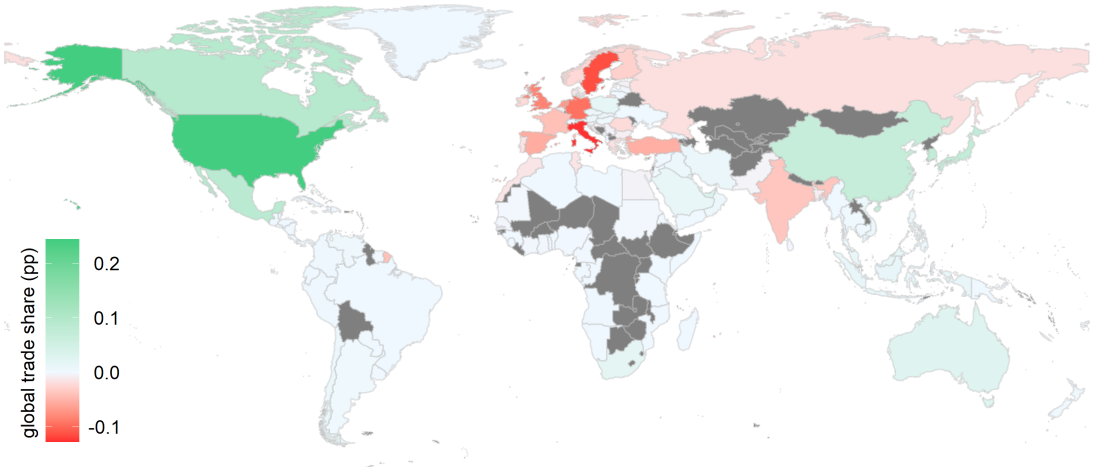

Note: Figure 9 shows the change in global export shares across countries in response to a 10% EU carbon tax on maritime transport. Countries in red reduce global export shares relative to the baseline, while countries in green expand their share of global exports.

competitors. Unilaterally taxing transport emissions in a carbon-intensive economy instead might therefore be more efficient from an environmental standpoint.

Emission savings cannot compensate for the loss in real wages: On aggregate, I find that a 10% EU carbon tax on container traffic decreases global welfare by 0.3% and saves 22 million tons of CO2. To compare these welfare components to each other, I translate emission savings into USD using estimates of the social cost of carbon (SCC). The SCC aims to capture the long-run economic damage caused by emitting an additional ton of CO2 today. As SCC estimates in the literature vary considerably, I present results for two different estimates that have recently gained a lot of attention. The Biden administration for example uses an SCC estimate of 51 USD/tCO2 to evaluate the environmental impact of prospective policies, while a recent study of the US Environmental Protection Agency puts the SCC at 190 USD/tCO2 (EPA, 2023).

Table 3 summarizes the combined effect of environmental and non-environmental welfare components. Across SCC estimates, the EU carbon tax decreases global welfare roughly 240 billion USD. The welfare loss related to the reduction in trade hereby exceeds the monetary gain of emission savings by an average factor of 42 to 89. Figure 10 and Table 14 in appendix E.6 complement these findings with welfare estimates for each individual country. Even after accounting for differences in the social cost of carbon across countries by using regional damage estimates of Nordhaus and Boyer (2000) and calculating

Table 3: Welfare effects of 10% carbon tax

change	Global	EU	non-EU
real wage (bn USD)	-244.41	-127.60	-116.45
transport emissions (mn ton CO2)	-18.14	-7.24	-10.56
production emissions (mn ton CO2)	-4.18	-2.10	-2.08
welfare SCC50 (bn USD)	-241.66	-125.50	-115.82
ratio SCC50	-88.98	-60.82	-184.20
welfare SCC190 (bn USD)	-238.54	-124.19	-114.05
ratio SCC190	-41.62	-37.47	-48.47

Note: The table presents welfare estimates of a 10% carbon tax on EU segments. The benchmark setting assumes no allocation barriers and an idle rate of 0.2. Social welfare is based on an assumed social cost of carbon (SCC) of 50 and 190 USD per ton of CO2. Ratios divide the absolute change in real wages by the sum of social carbon benefit (in USD) and tax revenue (in EU countries).

Figure 10: Social welfare response to 10% carbon tax

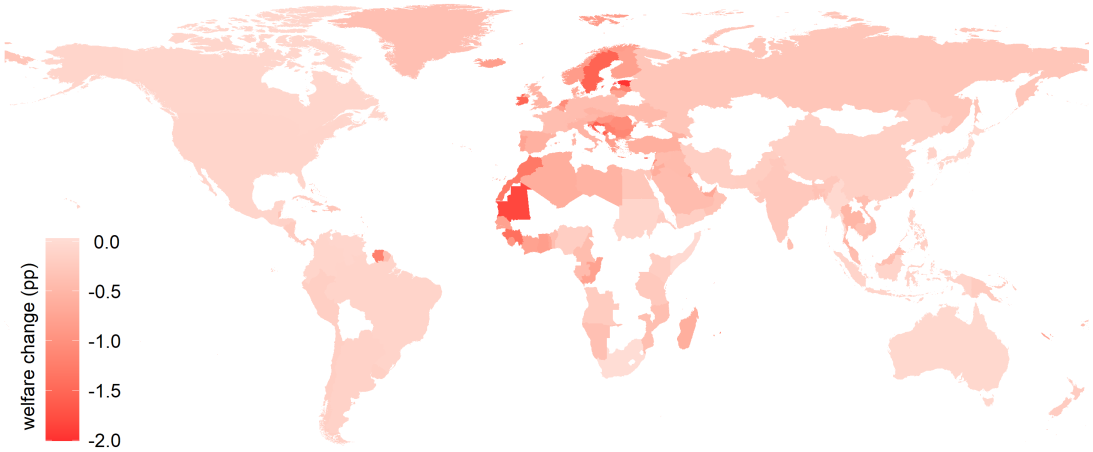

(a) relative change (in pp)

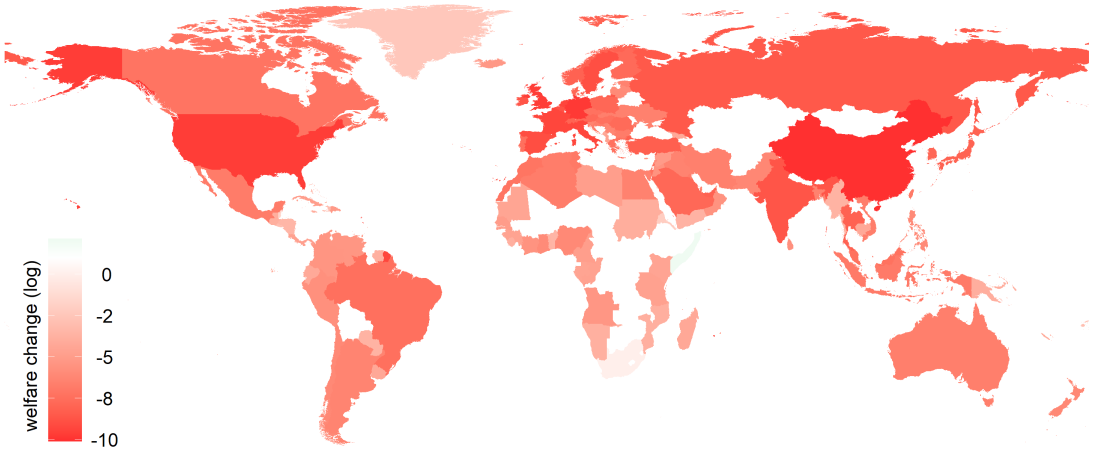

(b) absolute change (in billion USD)

country-specific damage weights as in Shapiro (2016), the welfare effects never become positive. Imposing a unilateral carbon tax therefore does not appear to be welfare enhancing, even after accounting for the environmental benefits of the policy.

7 Conclusion

Can unilateral policy decarbonize maritime trade? The results suggest that carbon taxes can indeed achieve large-scale emission savings and steer the sector towards its self-proclaimed emission target, but only at an enormous cost to welfare. Even in the most optimistic scenario, income losses from lower levels of trade exceed the environmental benefits of the policy by several orders of magnitude. This echoes similar findings in the

literature which has long emphasized that the gains from trade vastly outpace its environmental cost (Shapiro, 2016). Internalizing the environmental externalities of maritime trade via carbon taxes appears to be a very costly endeavour in the short run.

A lingering question in this regard is whether a temporary drop in welfare is acceptable as long as these policies accelerate the transition to a carbon-neutral transport technology. This paper does not attempt to resolve the intertemporal trade-off between the short and the long run. Nonetheless, it should serve as a cautionary tale against an overly optimistic picture in which the advent of some future technology grants absolution for any loss in the meantime. It may take decades until carbon-neutral vessels can be deployed at scale which means that the global fleet will continue to exhibit large levels of heterogeneity for many years to come. The cost incentives of shippers described in this paper will therefore continue to play an important role. Similarly, overcapacity will continue to diminish the operational efficiency of the fleet and raise global transport emissions, unless transport providers come up with new solutions to maintain high load factors, especially in the presence of negative demand shocks. As current order books forecast an unparalleled period of capacity expansion, issues of overcapacity could plague the profitability and environmental progress of the industry for many years to come.

At the same time, insights developed in this paper point to other potential avenues of sectoral decarbonization. First, unilateral carbon taxes in the shipping sector are more effective when imposed by countries with high emission intensity in production to maximize environmental spillovers. Second, directly incorporating environmental considerations into the design of shipping routes could limit demand dispersion and increase operational efficiency of the global fleet. Finally, policymakers could consider 'cash for clunkers' policies which subsidize the scrapping of inefficient vessels to lower overcapacity while raising average fleet efficiency. These approaches could deliver substantial emission savings while limiting the negative impacts of policy interventions and therefore represent promising avenues of future research.

References

Aichele, R. and G. Felbermayr (2015). Kyoto and Carbon Leakage: An Empirical Analysis of the Carbon Content of Bilateral Trade. *The Review of Economics and Statistics 97*(1), 104–115.

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

Alvarez, F. and R. E. Lucas (2007). General equilibrium analysis of the eaton–kortum model of international trade. *Journal of Monetary Economics 54*(6), 1726–1768.

Antweiler, W., B. R. Copeland, and M. S. Taylor (2001). Is free trade good for the environment? *American Economic Review 91*(4), 877–908.

Asturias, J. (2020). Endogenous transportation costs. *European Economic Review 123*, 103366.

Bombardini, M. and B. Li (2020, July). Trade, pollution and mortality in China. *Journal of International Economics 125*, 103321.

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

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

Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng (2023). Binscatter Regressions. *The Stata Journal*, 1–51.

Cherniwchan, J. (2017). Trade liberalization and the environment: Evidence from NAFTA and U.S. manufacturing. *Journal of International Economics 105*, 130–149.

Chung, S. (2014). Environmental regulation and foreign direct investment: Evidence from South Korea. *Journal of Development Economics 108*, 222–236.

Copeland, B. R., J. S. Shapiro, and M. Scott Taylor (2022). Chapter 2 - globalization and the environment. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics: International Trade, Volume 5*, pp. 61–146.

Copeland, B. R. and M. S. Taylor (1994). North-South Trade and the Environment. *The Quarterly Journal of Economics 109*(3), 755–787.

Coşar, A. K. and B. Demir (2018). Shipping inside the box: Containerization and trade. *Journal of International Economics 114*, 331–345.

Cristea, A., D. Hummels, L. Puzzello, and M. Avetisyan (2013). Trade and the greenhouse gas emissions from international freight transport. *Journal of Environmental Economics and Management 65*(1), 153–173.

Ducruet, C., R. Juhász, D. K. Nagy, and C. Steinwender (2024). All aboard: The effects of port development. *Journal of International Economics 151*, 103963.

Eaton, J. and S. Kortum (2002). Technology, Geography, and Trade. *Econometrica 70*(5), 1741–1779.

EPA (2023). EPA Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances. Technical Report EPA-HQ-OAR-2021-0317, US EPA.

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

Farrokhi, F. and A. Lashkaripour (2023). Can Trade Policy Mitigate Climate Change? *mimeo*.

Fowlie, M. and M. Reguant (2018). Challenges in the measurement of leakage risk. *AEA Papers and Proceedings 108*, 124–29.

Fuchs, S. and W. F. Wong (2022). Multimodal Transport Networks. *Federal Reserve Bank of Atlanta, Working Papers 2022*(13).

Ganapati, S., W. F. Wong, and O. Ziv (2024). Entrepôt: Hubs, scale, and trade costs. *American Economic Journal: Macroeconomics 16*(4), 239–78.

Gaulier, G. and S. Zignago (2010). Baci: International trade database at the product-level. the 1994-2007 version. Working Papers 2010-23, CEPII.

Hanna, R. (2010). Us environmental regulation and fdi: Evidence from a panel of us-based multinational firms. *American Economic Journal: Applied Economics 2*(3), 158–89.

Hansen-Lewis, J. and M. M. Marcus (2022). Uncharted waters: Effects of maritime emission regulation. Working Paper 30181, National Bureau of Economic Research.

Heiland, I., A. Moxnes, K. H. Ulltveit-Moe, and Y. Zi (2022). Trade From Space: Shipping Networks and The Global Implications of Local Shocks. *mimeo*.

IMO (2020). Fourth IMO GHG Study 2020.

Kortum, S. and D. A. Weisbach (2022). Optimal Unilateral Carbon Policy. *mimeo*.

Kucheryavyy, K., G. Lyn, and A. Rodríguez-Clare (2023). Grounded by gravity: A well-behaved trade model with industry-level economies of scale. *American Economic Journal: Macroeconomics 15*(2), 372–412.

Leibovici, F. and J. Dunn (2023). Navigating the Waves of Global Shipping: Drivers and Aggregate Implications. Technical Report 2023-002.

Lenzen, M., K. Kanemoto, D. Moran, and A. Geschke (2012). Mapping the structure of the world economy. *Environmental Science & Technology 46*(15), 8374–8381.

Lenzen, M., D. Moran, K. Kanemoto, and A. Geschke (2013). Building eora: A global multi-region input-output database at high country and sector resolution. *Economic Systems Research 25*(1), 20–49.

Lugovskyy, V., A. Skiba, and D. Terner (2024). Unintended Consequences of Environmental Regulation of Maritime Shipping: Carbon Leakage to Air Shipping. *mimeo*.

Monios, J. (2023). When smooth space becomes turbulent: The collapse of Hanjin Shipping and the immobilisation of ships, containers, goods and people. *Environment and Planning A: Economy and Space 55*(2), 320–338.

Nordhaus, W. (2015). Climate Clubs: Overcoming Free-riding in International Climate Policy. *American Economic Review 105*(4), 1339–1370.

Nordhaus, W. D. and J. Boyer (2000). *Warming the World: Economic Models of Global Warming*. The MIT Press.

Notteboom, T., A. Pallis, and J.-P. Rodrigue (2022). *Port Economics, Management and Policy* (1st ed.). London: Routledge.

Redding, S. J. and M. A. Turner (2015). Transportation Costs and the Spatial Organization of Economic Activity. In *Handbook of Regional and Urban Economics*, Volume 5, pp. 1339–1398. Elsevier.

Schott, P. K. (2014). The relative sophistication of Chinese exports. *Economic Policy 23*(53), 6–49.

Shapiro, J. S. (2016). Trade Costs, CO $_2$, and the Environment. *American Economic Journal: Economic Policy 8*(4), 220–254.

Shapiro, J. S. (2020). The Environmental Bias of Trade Policy*. *The Quarterly Journal of Economics 136*(2), 831–886.

Shapiro, J. S. and R. Walker (2018). Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade. *American Economic Review 108*(12), 3814–3854.

Simonovska, I. and M. E. Waugh (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics 92*(1), 34–50.

UNCTAD (2022). *Navigating stormy waters*. Number 2022 in Review of maritime transport, United Nations Conference on Trade and Development. Geneva: United Nations.

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

A Data sources

A.1 Fuel efficiency of the global container fleet

This section describes the main datasets and construction process of ship-level measures of fuel efficiency.

Ship movements: To start with, I use AIS port-of-call data from 2014-2018. Each time a ship is entering or leaving a port, a satellite records the time, location and draught of the vessel in question. I use this information to track vessel movements across ports and thereby reconstruct the flow of traffic across the transport network. The data contains around 470,000 port calls made by 5,100 container vessels at over 1100 ports each year.

Port-to-port distance: To compute the distance between ports, I use the recorded latitude and longitude of AIS port call data and apply Dijkstra's algorithm to find the shortest path between two adjacent port calls. The recovered sea distance represents a lower bound of the actual travel distance of the vessel which remains unobserved. Emission estimates based this distance measure may therefore underestimate the true environmental cost of trade.

Container traffic: The recorded draught at entry and exit of ports can be used to infer how much cargo a vessel is carrying at any moment in time. To do so, I compare the observed draught recorded by the satellite, H_A, with the minimum and maximum draught of the vessel, H_B and H_S, commonly referred to as ballast and scantling draught in the industry. Both are reported by the vessel manufacturer in a separate dataset which also includes the maximum carrying capacity of each vessel, known as the ship's deadweight tonnage (dwt). I follow Heiland et al. (2022) and assume that ballast draught $H_B = 0.55H_S$ which means that a ship is considered to travel empty if the recorded draught is at or below 55% of its scantling draught. The cargo weight of each individual trip can then be estimated using the following formula:

$$dwt_{current} = dwt_{max} * (H_A - H_B)/(H_S - H_B).$$

Fuel consumption: A separate dataset provided by S&P Global reports hourly fuel consumption estimates of container ships between 2018 and 2021. These engineering

estimates are based on the specific engine configuration installed on each ship and account for the vessel's current capacity utilization, speed and location. To capture fuel efficiency differences between ships, I compute the annual efficiency ratio (AER) of each vessel by dividing total fuel consumption by the ship's total carrying capacity (dwt) and annual travel distance recovered from port-of-call data. This efficiency ratio indicates how much fuel a vessel consumes per ton-km of supplied service capacity in 2018[65].

Data used to compute efficiency ratios exclusively focuses on fuel consumption at sea and therefore ignores any consumption created while manoeuvering or anchoring within ports. To ensure a sufficient amount of information for each vessel, I require each ship to be active for at least 720 hours at positive speed above ballast draught within a given year. In case of missing data in 2018, I use efficiency ratios of the same vessel from subsequent years. If vessel data is unavailable across years, I impute efficiency ratios by taking an average of fuel efficiency of vessels which share the same carrying capacity, year of build and engine power. This procedure allows me to recover the fuel efficiency of 5,333 container vessels of which 17% are imputed.

Transport emissions: When computing transport emissions of an individual journey, I multiply each vessel's efficiency ratio with the observed amount of transport work. As AERs reflect a vessel's fuel efficiency at the average load factor in 2018, I correct for diverging load factors on each journey using load factor elasticity β recovered in appendix D.1.

A.2 Identifying shipping routes

Container ships typically travel along a circular route which involves a fixed sequence of ports. This behaviour resembles buses in an urban transport network and differentiates them from other forms of maritime transport such as bulk shipping which are more comparable to taxis (Brancaccio et al., 2020).

To identify circular container routes in the data, I rely on AIS port-of-call data of container vessels between 2014 and 2018. In line with the model, I aggregate individual trips

[65]Note that a vessel's carrying capacity can also be expressed as the total number of containers the ship can carry. AERs can therefore be expressed in units of ton-km or TEU-km.

to the country level and collapse domestic journeys to a single entry. The result is a ship-level chain of country-to-country journeys over the entire sample period. Container routes can thus be thought of as reoccurring sequences within the chain which share a common start and end point.

To divide the chain into routes, I use the most frequently visited country within a given year as a starting point and define a route as the immediate sequence of country stops before returning to the initial country. This procedure yields around 9,000 circular routes per year which account for 99% of global container traffic. Further details are provided in Table 4 which lists the five most and least frequented routes in 2018.

Table 4: Shipping routes in 2018

route	volume (1000 ton)	rotations
top 5 routes		
CHN-KOR-CHN	377960	6483
CHN-HKG-CHN	195246	3206
CHN-TWN-CHN	187134	4506
CHN-JPN-CHN	179712	7198
BRA-URY-ARG-URY-BRA	138907	985
bottom 5 routes		
CHN-MAC-CHN	4	2
ROU-TUR-BGR-ROU	3	3
IRN-QAT-ARE-IRN	3	3
TUR-RUS-GRC-TUR	3	3
IRN-ARE-QAT-IRN	2	3

The table shows the five most and least frequent shipping routes in the network based on transport volume. Routes are identified following the procedure outlined in appendix A.2.

A.3 Additional data for counterfactual analysis

To estimate policy counterfactuals, I complement data from section 2 with two additional data sets.

Trade data: Bilateral trade data is obtained from the BACI database, maintained by the Centre d'etudes Prospectives et d'Informations Internationales (CEPII). It contains data on bilateral trade flows of 227 countries in 2018. To aggrgate BACI data to containerized and non-containerized sectors, I follow the procedure described in appendix D.3.

Global input-output data: The second new dataset is the EORA Global Supply Chain

Database which is a multi-region input-output table and contains information on input-output linkages, country-level production and consumption and emission intensities for 190 countries. I rely on the simplified Eora26 table which consists of 26 sectors and is publicly available at https://www.worldmrio.com/eora26/.

Table 5 describes how to map the 26 sectors of the original data into the three sectors of the model. To compute Cobb-Douglas expenditure shares φ and value added shares γ for each sector, I follow the procedure explained in Caliendo and Parro (2015).

Production emission intensity: EORA also reports total CO2 emissions created by each sector from the production of intermediary and final goods. I use this information to compute emission intensities by dividing total emissions by total output in each country-sector. Emission intensity is then measured in units of ton of CO2 per USD of output.

Table 5: Mapping EORA sectors to the model

Model sector	EORA sector
containerized trade	Fishing, Food and Beverages, Textiles and Wearing Apparel, Wood and Paper, Metal Products, Electrical and Machinery, Transport Equipment, Other Manufacturing, Recycling, Reexport and Reimport, Others
non-containerized trade	Agriculture, Mining and Quarrying, Petroleum Chemical and Non-Metallic, Electricity Gas and Water, Financial Intermediation and Business Activities
non-tradable	Construction, Maintenance and Repair, Wholesale Trade, Retail Trade, Hotels and Restraurants, Transport, Post and Telecommunications, Public Administration, Education Health and Other Services, Private Households

B Additional graphs

B.1 Change in emission determinants over time

Figure 11: Global transport emissions of container ships (2016-2018)

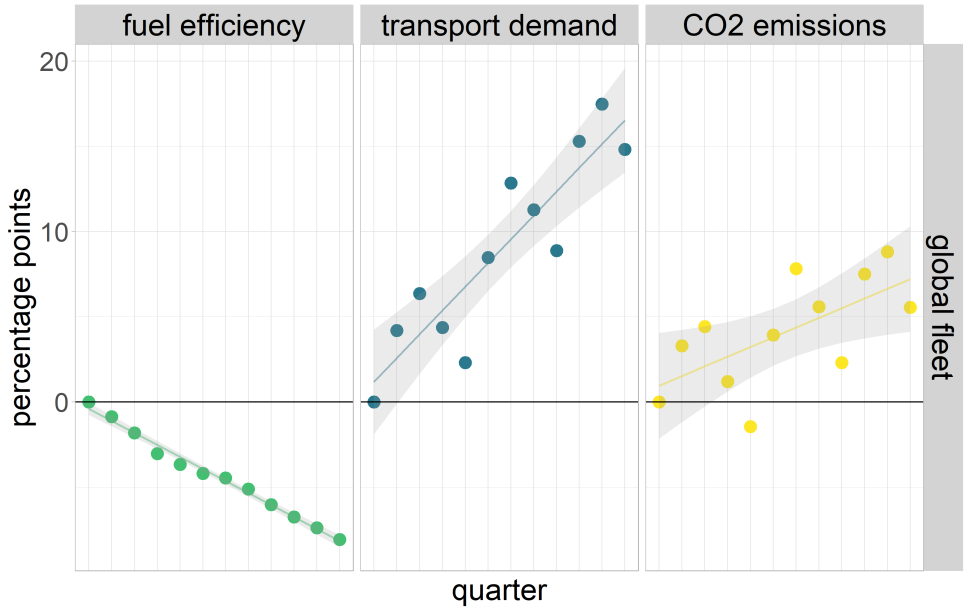

Note: Figure 11 displays the changes in fuel efficiency, transport demand and CO2 emission relative to the first quarter of 2016.

B.2 Load factors and network traffic

Figure 12: Load factors and segment traffic by ship size

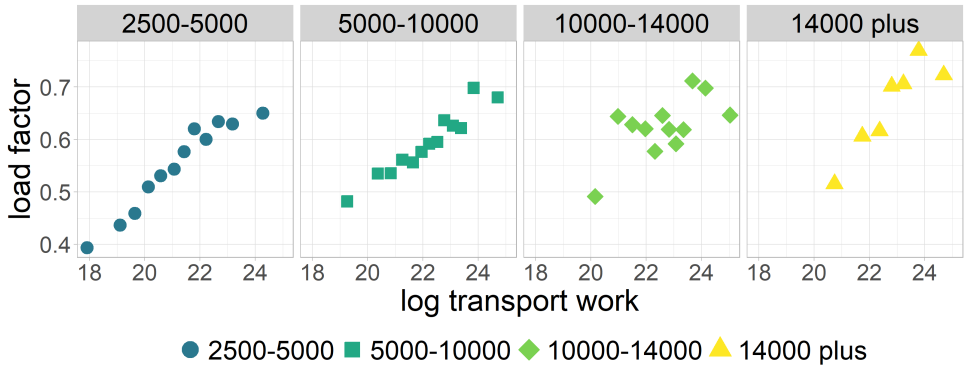

Note: This Figure plots the relationship between average load factors and segment traffic for different vessel size bins. It is otherwise identical to Figure 2 in the main text.

B.3 Fleet characteristics

Figure 13: Supply of transport capacity

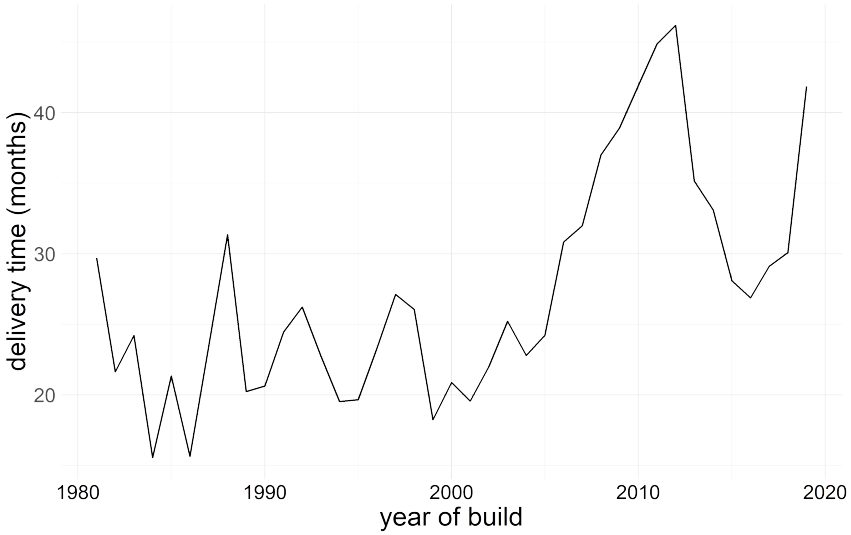

(a) delivery times of container ships

(b) global fleet in use

Note: The Figure grants insights into the supply of containerized transport capacity. Panel a) shows average delivery times (in months) of new container ships between 1990 and 2020. Panel b) displays the daily share of global transport capacity in use in 2018. For this graph, a ship is considered 'idle' if it has travelled for less than 10km per day for at least 7 consecutive days.

C Additional derivations

C.1 Sourcing probability

To derive the sourcing probability in equation 4, I closely follow Eaton and Kortum (2002). Using the stochastic transport cost from equation 3, the price of a good produced in country i which is shipped along route r before being sold in country j is

$$p_{ijnr} = \frac{c_{in}\kappa_{ijn}\tilde{\tau}_{ijr}}{\epsilon_{ijnr}}$$

where ϵ_{ijnr} is an ideosyncratic route shock drawn from a Fréchet distribution. Random variation in route costs renders the choice of optimal supplier-route combinations non-deterministic. The probability that price p_{ijnr} represents the lowest price for consumers in country j can be written as

$$\pi_{ijnr}(\omega) \equiv \Pr\left\{p_{ijnr} \leq \min_{i' \in I \setminus i, r' \in R_{ij} \setminus r} p_{i'jnr'}\right\} = \int \prod_{i',r'} (1 - G_{i'jnr'}) dG_{ijnr} dp \tag{20}$$

where

$$G_{ijnr} = \Pr\left\{p_{ijnr} < p\right\}$$
$$\prod_{i',r'} (1 - G_{i'jnr'}) = \prod_{i',r'} \Pr\left\{p_{i'jnr'} \leq p\right\}.$$

Using the cumulative distribution function of route shocks ϵ_{ijnr} specified in equation 2, we get

$$G_{ijnr} = 1 - \exp\left[-p^{\theta}(c_{in}\kappa_{ijn}\tilde{\tau}_{ijr})^{-\theta}\right]$$
$$dG_{ijnr} = -\exp\left[-p^{\theta}(c_{in}\kappa_{ijn}\tilde{\tau}_{ijr})^{-\theta}\right] \theta p^{\theta-1}(c_{in}\kappa_{ijn}\tilde{\tau}_{ijr})^{-\theta}$$
$$\prod_{i',r'} (1 - G_{i'jnr'}) = \exp\left[-p^{\theta} \sum_{i'} \left((c_{i'n}\kappa_{i'jn})^{-\theta} \cdot \sum_{r'} \tilde{\tau}_{i'jr'})^{-\theta} \right)\right].$$

Plugging these three expressions into equation 20 and integrating over all values of p, we get the sourcing probability shown in equation 4 the main text:

$$\pi_{ijnr} = \frac{\left[c_{in}\kappa_{ijn}\tilde{\tau}_{ijr}\right]^{-\theta}}{\sum_{i' \in I} \left[(c_{i'n}\kappa_{i'jn})^{-\theta} \cdot \sum_{r' \in R_{i'j}} \tilde{\tau}_{i'jr'}^{-\theta} \right]}$$

C.2 Optimal allocation of surplus capacity

In this section, I derive the optimal allocation of surplus capacity in a setting where clean ships have been allocated according to equation 12 and dirty capacity, after idling as many vessels as possible, $(1-\eta)B_d$ exceeds the remaining transport demand B'_d.

In the absence of surplus capacity, each route receives the minimum amount of capacity required to serve transport demand. As explained in the main text, the route-segment with the highest demand determines the number of route rotations $\frac{\Xi_{L,max}}{N_s}$ resulting in a minimum route capacity of $\frac{\Xi_{L,max}}{N_s} N_s \sum_h d_{L,h} = \Xi_{L,max} \sum_h d_{L,h}$. In presence of surplus capacity, vessels no longer operate at full capacity on route-segment (L, max). Assuming demand is equally shared among all vessels operating on the same route, this increases the number of route rotations to $\frac{\Xi_{L,max}}{\mu_{L,max} N_s}$ where $\mu_{L,max} < 1$ represents the load factor on the top segment. For a given level of $\mu_{L,max}$, total capacity supplied to route L in presence of surplus capacity is thus

$$\text{total capacity with surplus} = \frac{\Xi_{L,max}}{\mu_{L,max}} \sum_h d_{L,h}. \tag{21}$$

To facilitate the discussion, I assume a common maximum load factor $\mu_{L,max}$ across all routes that receive surplus capacity. Parameter $\mu_{L,max}$ governs the amount of surplus capacity each route can accommodate. A common maximum load factor means that routes differ in the *level* of surplus capacity they receive but all share the same *relative* capacity increase compared to minimum capacity $\Xi_{L,max} \sum_h d_{L,h}$. This keeps the allocation of surplus capacity tractable and can be rationalized by a common minimum fuel efficiency threshold for continued route operation that is shared by all shippers.

Surplus capacity increases the number of full and empty trips on a route-segment to $\frac{\Xi_{L,h}}{\mu_{L,max} N_s}$ and $\frac{\Xi_{L,max} - \Xi_{L,h}}{\mu_{L,max} N_s}$ respectively, leading to a fuel expenditure on route-segments of

$$F_{L,h} = \left(\hat{\delta}_{L,s}^{\mu=\mu_{L,max}} \cdot \frac{\Xi_{L,h}}{\mu_{L,max} N_s} + \hat{\delta}_{L,s}^{\mu=0} \cdot \frac{\Xi_{L,max} - \Xi_{L,h}}{\mu_{L,max} N_s} \right) \cdot d_{L,h}$$

and total route cost to

$$F_L = \frac{1}{\mu_{L,max}} \left(\delta_{L,s}^{\mu=\mu_{L,max}} \cdot \sum_h \Xi_{L,h} d_{L,h} + \delta_{L,s}^{\mu=0} \cdot \sum_h (\Xi_{L,max} - \Xi_{L,h}) d_{L,h} \right). \tag{22}$$

To compare the impact of surplus capacity across routes I derive two objects. First, the increase in route cost from a marginal decrease in load factor $\mu_{L,max}$. Second, the increase in surplus capacity on route L associated with the same decrease in load factor. Both objects are required because a marginal decrease in $\mu_{L,max}$ causes different levels of surplus capacity across routes. Dividing the former by the latter therefore harmonizes the cost increase and allows me to compare the detrimental impact of surplus capacity across routes.

Surplus capacity on route L is equal to the difference between total capacity inclusive of surplus and the minimum capacity required to serve transport demand:

$$B_{surplus,L} = \frac{1-\mu_{L,max}}{\mu_{L,max}} \cdot \sum_h \Xi_{L,max} d_{L,h}.$$

Decreasing load factor $\mu_{L,max}$ results in an increase in surplus equal to

$$-\frac{\mathrm{d}B_{surplus,L}}{\mathrm{d}\mu_{L,max}} = \frac{1}{(\mu_{L,max})^2} \cdot \sum_h \Xi_{L,max} d_{L,h}. \tag{23}$$

The corresponding marginal change in route cost F_L in equation 22 is

$$-\frac{\mathrm{d}F_L}{\mathrm{d}\mu_{L,max}} = \frac{1}{(\mu_{L,max})^2} \cdot F_L \; - \; \frac{\mathrm{d}\delta_{L,s}^{\mu=\mu_{L,max}}}{\mathrm{d}\mu_{L,max}} \cdot \frac{1}{\mu_{L,max}} \cdot \sum_h \Xi_{L,h} d_{L,h}. \tag{24}$$

Plugging equation 22 into equation 24 and dividing equation 24 by equation 23, we get

$$\begin{aligned} C_L &= -\frac{\mathrm{d}F_L}{\mathrm{d}\mu_{L,max}} \quad \cdot \quad -\frac{\mathrm{d}\mu_{L,max}}{\mathrm{d}B_{surplus,L}} \\ &= -\frac{\mathrm{d}F_L}{\mathrm{d}\mu_{L,max}} \quad \cdot \quad \frac{(\mu_{L,max})^2}{\sum_h \Xi_{L,max} d_{L,h}} \\ &= \left(\delta_{L,d}^{\mu=\mu_{L,max}} - \frac{\mathrm{d}\delta_{L,s}^{\mu=\mu_{L,max}}}{\mathrm{d}\mu_{L,max}} \mu_{L,max} \right) \cdot \bar{\mu}_L \quad + \quad \delta_{L,d}^{\mu=0} \cdot (1-\bar{\mu}_L) \end{aligned}$$

which is identical to equation 14 in the main text.

C.3 Solving the general equilibrium model

To close the model in section 4, I adopt the general equilibrium framework of Caliendo and Parro (2015). The framework features two types of goods. Intermediate goods ω_n from industry n are used as inputs in the production of composite goods m_{in} in country i. Composite goods in turn either serve as final goods for consumers or become inputs in the production of intermediate goods. Intermediate and composite goods therefore create a system of roundabout production. While both serve as inputs for each other, intermediate goods are only produced from domestic composite goods while composite goods source intermediate inputs globally from the lowest cost suppliers as long as intermediates ω_n are tradable.

I aggregate industries into three sectors ($N = 3$): two tradable sectors (ct = containerized and nct = non-containerized) and a non-tradable sector nt such that $n \in [ct, nct, nt]$. Composite goods therefore have to source inputs ω_{nt} domestically but can source ω_{ct} and ω_{nct} globally. While all three sectors transport goods through a sector-specific cost matrix A_n, I only consider cost changes in containerized sector ct.

Consumption: Consumers in each country i choose from a variety of composite goods m_{in} to maximize Cobb-Douglas utility

$$U_i = \prod_{n}^{N} m_{in}^{\varphi_n} \quad \text{where} \quad \sum \varphi_n = 1$$

where φ_n represents the Cobb-Douglas industry shares.

Intermediate goods production: Firms in country i produce a continuum of intermediate goods in each sector n using labour l_{in} and composite goods m_{in} as inputs. Countries differ in their efficiency of production which is represented by Ricardian technology parameter z_{in}. The sector-specific production function for intermediate good ω is thus

$$q_{in}(\omega) = z_{in}[l_{in}]^{\gamma_{in}} \prod_{n'}^{N} \left[m_{in}^{n'}\right]^{\gamma_{in}^{n'}}$$

where $\gamma_{in} = 1 - \sum_{n'}^{N} \gamma_{in}^{n'}$ captures the share of value added of industry n in country i. Under perfect competition, the marginal cost of producers then becomes

$$c_{in} \equiv \frac{\Upsilon_{in} \mathrm{w}_i^{\gamma_{in}} \prod_{n'}^{N} P_{in'}^{\gamma_{in}^{n'}}}{z_{in}}$$

where w_i represents the wage in country i, $P_{in'}$ is the price of composite good $m_{in'}$ from sector n' and $\Upsilon_{in} = \prod_{n'}^{N} \left(\gamma_{in}^{n'}\right)^{\gamma_{in}^{n'}} \left(\gamma_{in}\right)^{\gamma_{in}}$ is a constant.

Composite goods production: Composite goods are produced as a CES aggregate of intermediate goods which are indexed in each sector by $\omega_n \in \Omega_n$. In non-tradable industries nt, intermediate inputs are sourced domestically at marginal cost c_{in}. In tradable sectors ct and nct, intermediate goods are sourced internationally from the lowest-cost supplier after taking into account trade costs κ_{ijn} such as tariffs and transport costs τ_{ijn} which are determined by the optimal allocation of ships as shown in section 4.2. The expected price of composite good m_{jn} in sector n of country j is then

$$P_{jn} = A_n \left[\sum_{i=1}^{I} c_i^{-\theta_n} \kappa_{ijn}^{-\theta_n} \tau_{ijn}^{-\theta_n} \right]$$

where A_n is a sector-specific constant.

Equilibrium in changes: In that follows, the hat-notation symbolizes the change in underlying variable $\hat{v} = \frac{v'}{v}$ where v' represents counterfactual values of v. An exogenous shock to transport costs on network segment (k, l), $t_{k,l}$, leads to a change in the marginal cost of production of intermediate goods in country i according to

$$\hat{c}_{in} = \hat{\mathrm{w}}_i^{\gamma_{in}} \prod_{k=1}^{N} P_{ik}^{\gamma_{ink}}. \tag{25}$$

This cost increase changes prices of composite goods in i to

$$\hat{P}_{in} = \left[\sum_{i=1}^{J} \pi_{ijn} [\hat{\tau}_{ijn} \hat{c}_{in}]^{-\theta_n} \right]^{-1/\theta_n}. \tag{26}$$

Bilateral trade shares between countries i and j respond to changes in production cost $\hat{c}_{in}$, transport cost $\hat{\tau}_{ijn}$ and price of the composite good $\hat{P}_{in}$ according to

$$\hat{\pi}_{ijn} = \left[\frac{\hat{c}_{in} \hat{\tau}_{ijn}}{\hat{P}_{in}} \right]^{-\theta_n}. \tag{27}$$

Total exports of country i adjust accordingly to

$$X_{in}^{'} = \sum_{k=1}^{N} \gamma_{ink} \sum_{j=1}^{I} \frac{\pi_{ijn}^{'}}{1+\kappa_{ijn}} X_{jk}^{'} + \alpha_{in} I_{i}^{'} \tag{28}$$

where $I_{i}^{'} = \hat{\mathrm{w}}_i \mathrm{w}_i L_i + \sum_{n=1}^{N} \sum_{i=1}^{I} \tau_{ijn}^{'} \frac{\pi_{ijn}^{'}}{1+\kappa_{ijn}} X_{in}^{'} + D_i$. To balance trade in the counterfactual (up to deficit shifter D_i), it must hold that

$$\sum_{n=1}^{N} \sum_{i=1}^{I} \frac{\pi_{ijn}^{'}}{1+\kappa_{ijn}} X_{in} - D_i = \sum_{n=1}^{N} \sum_{i=1}^{I} \frac{\pi_{jin}^{'}}{1+\kappa_{jin}} X_{jn}. \tag{29}$$

To solve this system of equations, I follow the solution algorithm of Caliendo and Parro (2015) which iteratively determines the change nominal wage vector w for which trade in equation 29 is balanced.

D Estimation details

D.1 Within-ship load factor elasticity

This section explains how to estimate within-ship load factor elasticity β using fuel consumption data explained in detail section A. The load factor elasticity captures the percentage change in fuel consumption per container-slot within vessels when increasing the vessel's load by one percentage point. Higher loads are generally associated with lower unit fuel consumption because the fixed cost of operating the vessel can be spread over a larger number of containers.

To estimate the load factor elasticity, I rely on hourly fuel consumption data of container vessels in 2018. To rule out inconsistencies in vessel operation, I focus on fuel consumption recorded at days with 24 hours of sail activity. The dependent variable is hourly unit fuel consumption per transport volume (measured in TEU) or per transport work (measured in TEU-km). The main variable of interest, load factors, are based on the average capacity utilization of the vessel on each day which directly observed by the satellite. In line the literature[66], I assume that 45% of total vessel capacity is available for cargo and compute load factors by rescaling the capacity utilization levels between 55% and 100% to the unit interval.

To control for unobserved vessel-specific characteristics which might affect the ship's unit fuel consumption such as different levels of maintenance or repair, I include ship-level fixed effects. Speed limits or other region-specific unobservables are controlled for via another set of sailing region fixed effects. The baseline estimation equation is thus

$$\log(\bar{F}_{s,r,t}) \;=\; \beta \cdot \mu_{s,t} \;+\; FE_s \;+\; FE_r \;+\; \varepsilon_{s,r,t} \tag{30}$$

where indices (s, r, t) refer to observations at the ship-region-day level. The main coefficient of interest is load factor elasticity β which captures the percentage change in unit fuel consumption from a one percentage point (pp) increase in load factor μ.

[66] Heiland et al. (2022) for example assume that a ship is sailing in ballast, meaning it does not carry any containers, until the observed draught exceeds 55% of the vessel's maximum (= scantling) draught.

Benchmark results shown in Table 6 indicate that the unit fuel consumption falls between 51% to 71% when raising load factors from 0% to 100%. Robustness checks presented in Table 7 confirm these estimates.

Table 6: Load factor elasticity within ships

Dependent var:	fuel per TEU	fuel per TEU-km	fuel per TEU-km	fuel per TEU-km
Model:	all	all	small ships	big ships
Variables				
load factor (pp)	-0.0057***	-0.0068***	-0.0071***	-0.0051***
	(0.0003)	(0.0002)	(0.0002)	(0.0003)
ship FE	yes	yes	yes	yes
sailing region FE	yes	yes	yes	yes
R^2	0.530	0.757	0.739	0.603
Observations	195,902	195,902	165,607	30,295

The table shows estimates of within-ship load factor semi-elasticity obtained from estimating equation 30 in appendix D.1 via OLS. Fuel efficiency is measured in units of transport volume in column 1 and transport work otherwise. Load factors are measured in percentage points (pp). Standard errors in parentheses are clustered at the sailing region level. Significance codes: ***: 0.01, **: 0.05, *: 0.1.

Table 7: Robustness: Load factor elasticity within ships

Model:	benchmark	load factor 50+	ECA	TEU>2500	seasonality
Variables					
load factor (pp)	-0.0068***	-0.0051***	-0.0070***	-0.0068***	-0.0069***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
ship FE	yes	yes	yes	yes	yes
sailing region FE	yes	yes	yes	yes	yes
ECA FE			yes		
region-day FE					yes
R^2	0.757	0.743	0.758	0.676	0.804
Observations	195,902	152,136	196,761	162,598	196,761

The table shows estimates of within-ship load factor semi-elasticity obtained from estimating equation 30 in appendix D.1 via OLS. Fuel efficiency is measured in units of transport work. Load factors are measured in percentage points (pp). Column 1 is identical to benchmark estimates presented in Table 6. Column 2 and 4 restrict the sample to ships travelling at a minimum load factor of 50% and ships with a minimum capacity of 2,500 TEU respectively. Columns 3 and 5 add emission control area (ECA) and daily weather changes via region-day fixed effects. Standard errors in parentheses are clustered at the sailing region level. Significance codes: ***: 0.01, **: 0.05, *: 0.1.

D.2 Between-ship fuel efficiency

Figure 14: Estimation of between-ship parameters δ_c and δ_d

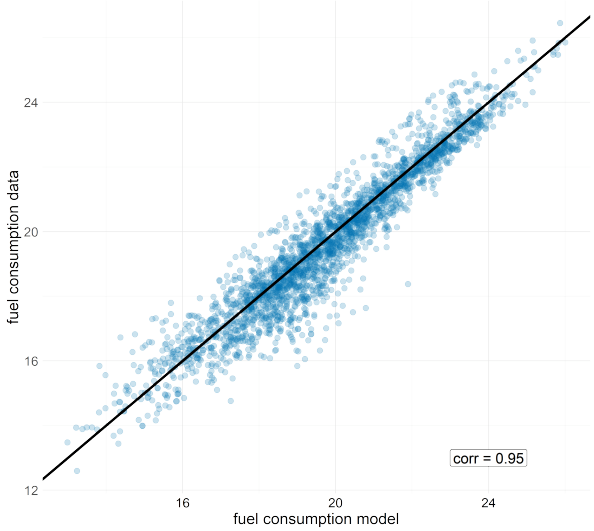

(a) $\ln F^{\text{data}}$ and $\ln F^{\text{model}}$

Coefficient	Estimate
$\delta_c^{\mu=0}$	0.932
$\delta_d^{\mu=0}$	1.864
$\text{corr}(\ln F^{\text{data}}, \ln F^{\text{model}})$	0.946

(b) Fuel per TEU-km

This Figure presents GMM estimates of equation 18 presented in section 5.1. Figure 14a illustrates the correlation (in logs) between observed and model-predicted fuel expenditure using optimal fuel efficiency parameters δc and δ_d which are listed in Figure 14b. To ensure no overlap in efficiency range of clean and dirty vessels, I impose $\delta_d^{\mu=1} \geq \delta_c^{\mu=0}$ when estimating parameters δ. GMM estimates assume a clean ship size $N_c = 10,000$ TEU, a dirty ship size $N_d = 1,000$ TEU and exogenous surplus parameter $\mu_{max} = 0.75$.

D.3 How to recover global transport cost

This section provides further details about how to recover transport cost matrix A following the procedure developed by Ganapati et al. (2024) (GWZ) and summarized in the main text.

To start with, I create a mapping between HS4 products and their dominant mode of transport. To do so, I use US customs data from 2018 maintained by Schott (2014) and available at https://sompks4.github.io/sub_data.html which breaks down import volumes by mode of transport. I label all HS codes as containerized if at least 80% of total import volume is transported onboard of container ships. I then apply the mapping to the publicly available BACI database (Gaulier and Zignago, 2010) which is maintained by Centre d'etudes Prospectives et d'Informations Internationales (CEPII) and contains data on bilateral trade flows of 227 countries in 2018. Aggregating trade according to the mapping then provides me with containerized trade flows at the origin-destination (i, j) level in 2018.

To recover transport cost of containerized trade, I assume that leg-level transport cost is identical within set of industries $n \in \bar{N}$ where $\bar{N} = \{\text{containerized, non-containerized}\}$. This changes equation 6 to

$$\Xi_{kl\bar{N}} = \sum_i \sum_j X_{ij\bar{N}} \left(\frac{\tau_{ij\bar{N}}}{\tau_{ik\bar{N}} t_{kl\bar{N}} \tau_{lj\bar{N}}} \right)^{\theta} \tag{31}$$

where elements $\left(\frac{\tau_{ij\bar{N}}}{\tau_{ik\bar{N}} t_{kl\bar{N}} \tau_{lj\bar{N}}} \right)$ are all based on cost matrix $A_{\bar{N}}$.

As leg-level container traffic $\Xi_{kl\bar{N}}$ and origin-destination level trade flows $X_{ij\bar{N}}$ are directly observed by AIS satellite data and aggregated BACI data described above, we can invert this system to recover cost matrix $A_{\bar{N}}$. The remaining challenge in doing so is that observed AIS traffic data only captures maritime traffic. It therefore may underestimate total containerized traffic on network legs (k, l) which are directly or indirectly connected via land borders.

To overcome this challenge, GWZ suggest a prediction method that does not require knowledge of the entire matrix of traffic flows. Instead of recovering transport cost from the full system of equations which requires information on $\Xi_{kl\bar{N}}$ for every segment, cost

elements of transport matrix $A_{\bar{N}}$ are predicted via the mapping

$$\hat{t}_{kl\bar{N}}^{-\theta} \quad = \quad \frac{1}{1+\exp(\mathbf{Y}\rho)} \quad \in \quad [0,1] \tag{32}$$

where matrix elements $\mathbf{Y}\rho$ are fitted values from an auxiliary regression that resembles a gravity equation

$$\begin{aligned}\mathbf{Y}\rho = \rho_0 + \rho_1 \cdot \ln \mathrm{d}_{kl} + \rho_2 \cdot \ln \Xi_{kl} + \rho_3 \cdot \ln \Xi_k + \rho_4 \cdot \ln \Xi_l \\ + \rho_5 \cdot 1_{backhaul} + \rho_6 \cdot 1\{k,l \in \text{land border}\} + \rho_7 \cdot 1\{k,l \in \text{landlocked}\}\end{aligned} \tag{33}$$

where ρ_0 is an intercept, ρ_1 captures the average sea distance between ports in countries k and l, ρ_2, ρ_3 and ρ_4 refer to total, incoming and outgoing traffic, ρ_5 indicates whether traffic in direction (k,l) is subject to backhaul problems due to higher demand in the mirror leg (l,k) (Wong, 2022), and ρ_6 and ρ_7 are dummy variables indicating whether countries share a land border or are landlocked.

Auxiliary equation 33 allows us to predict transport costs for any segment in the network for which independent variables captured by matrix $\mathbf{Y}$ are available. Feeding the fitted values of the auxiliary equation into mapping 32 ensures that all cost elements remain within the unit interval as required by the theoretical framework in section 4. While this procedure can be used to recover the entire cost matrix A, coefficients ρ remain unbiased when estimating them from a subset of legs (k,l) on which all traffic is maritime. This allows me to recover a transport cost matrix that rationalizes traffic on all legs, even though only a subset of them is used in the estimation.

To find cost elements $\hat{t}_{kl\bar{N}}^{-\theta}$ that can rationalize observed traffic and trade, coefficients ρ are estimated by minimizing equation 19 in the main text. GMM results are presented in Table 8. It should be noted that the sole purpose of coefficients ρ is to predict cost matrix A. They should not be interpreted by themselves and do not represent any causal relationship between independent variables and fitted values.

Table 8: prediction inputs for transport cost matrix

Coefficient		Estimate
ρ_0	(intercept)	10.64
ρ_1	(log distance)	-0.26
ρ_2	(log route traffic)	-1.38
ρ_3	(log outgoing traffic)	0.81
ρ_4	(log incoming traffic)	0.46
ρ_5	(back-haul)	-0.16
ρ_6	(land borders)	-1.84
ρ_7	(landlocked)	-6.40

The table shows GMM estimates of auxiliary regression 33 in appendix D.3. Estimated coefficients are not causal and only used to find transport cost matrix A that can rationalize observed leg-level traffic Ξ_{kl} and origin-destination trade flows X_{ij} in equation 6.

D.4 Decomposition of scale effects

This section illustrates how the theoretical framework can be used to obtain theory-consistent scale effect estimates that can be directly compared to reduced-form evidence in the literature. Prior attempts to estimate scale effects in the maritime transport network at the leg-level do not provide a microfoundation and instead pursue an approach that is agnostic to the underlying mechanism. Ganapati et al. (2024) (GWZ) for example specify the following equation to estimate scale effects in the shipping industry

$$\ln(t_h^\theta - 1) = \lambda_0 + \lambda_1 \cdot \ln \Xi_h + \lambda_2 \cdot \ln d_h + \varepsilon_h \tag{34}$$

where λ_1 can be interpreted as a leg-level scale elasticity between transport cost $(t_h^\theta - 1)$ and volume (Ξ_h) after deflating estimates by trade elasticity θ. To interpret λ_1 as a causal estimate, a geography-based instrument[67] is used to control for endogenous changes in transport demand Ξ_h due to unobserved cost components. For a commonly used trade elasticity of $\theta = 4$ (Simonovska and Waugh, 2014), the authors report a leg-level elasticity

[67]The instrument exploits the fact that segments h which are located on the shortest path between origin i and destination j should naturally receive higher levels of traffic. Summing across all origins and destinations then provides an index for each segment's natural suitability for global containerized trade. As this index is purely based on geography, it acts as an exogenous demand shifter. For the construction of the instrument, I refer to equation 12 of the original paper.

of -0.06 which for an average route-length of 3 translates to an origin-destination scale elasticity of -0.17.

Conversely, scale effects in this paper arise endogenously from the optimal allocation of transport capacity. To compare the empirical estimates of the literature to the model-implied scale elasticity of the theoretical framework, I proceed in three steps. In a first step, I follow the reduced-form approach of GWZ and estimate scale effect λ_1 using transport cost recovered in section 5.2 and observed levels of traffic in 2018. Next, I use the framework to derive a theory-consistent estimation equation and take it to the data. Third, I estimate scale effects directly from the calibrated model.

Scale effects in maritime shipping generally describe a decrease in transport cost for higher levels of traffic. While the empirical approach of GWZ assumes a direct relationship between traffic and cost (λ_1), the theoretical framework instead decomposes the overall effect into two elasticities: a fuel elasticity α_1 capturing the relationship between unit fuel consumption and traffic (= stylized fact 3), and a pass-through elasticity σ_1 governing the relationship between fuel efficiency and transport cost. The overall scale effect can then be computed as the product of both elasticities ($\lambda_1 = \alpha_1 \cdot \sigma_1$). To conduct a theory-consistent estimation of scale effects, I therefore need to derive an estimation equation for each elasticity.

Fuel and pass-through elasticities can either be estimated from external data or fully rely on the model-generated data based on the optimal allocation of ships. Each approach will result in a slightly different estimation equation.

When using external data, fuel elasticity α_1 describes the relationship between fuel efficiency $\bar{F}_h$ and transport volume Ξ_h. This allows me to use the same geography-based instrument as GWZ to control for unobserved cost shifters that simultaneously affect transport volume and unit fuel consumption. Elasticity α_1 can then be interpreted as a causal effect describing how fuel efficiency of vessels on segment h respond to a change in transport volume. To derive a theory-consistent estimation equation for fuel elasticity that allows for this instrumentation strategy, I simplify the model by assuming that a segment h that is part of multiple routes L is always served by the same type of vessel

($N_{L,h,s} = N_{h,s}$), receives a constant share of total segment volume ($\frac{\Xi_{L,h}}{\Xi_h} = w$) and is not subject to demand dispersion ($\Xi_{L,h} = \Xi_{L,max}$). This reduces the fuel expenditure shown in equation 15 to

$$F^*_{L,h} = \delta^{\mu=1} \cdot \Xi_{L,h} d_{L,h} \cdot n$$

where $n = \frac{\Xi_{L,h}}{N_{h,s}}$ is the number of segment rotations on route L which is identical across routes under the simplifying assumptions. Summing across routes L, we can express total fuel expenditure on segment h as

$$F^*_h = \delta^{\mu=1} \cdot n \cdot d_h \cdot \sum_L \Xi_{L,h}$$

$$F^*_h = \delta^{\mu=1} \cdot d_h^{\text{travel}} \cdot \Xi_h$$

where d_h^{travel} represents the total distance (in km) travelled on segment h and Ξ_h is the total number of containers shipped on segment h. Observed fuel efficiency $\bar{F}_h^{\text{data}}$, defined as total fuel expenditure F_h divided by total transport work W_h^{data}, can thus be decomposed into the following estimation equation

$$\ln \bar{F}_h^{\text{data}} \;=\; \alpha_0 \;+\; \alpha_1^{\text{data}} \cdot \ln \Xi_h^{\text{data}} \;+\; \alpha_2 \cdot \ln \frac{d_h^{\text{travel}}}{W_h^{\text{data}}} \;+\; \upsilon_h \tag{35}$$

where α_1^{data} captures the fuel elasticity of observed fuel efficiency $\bar{F}_h^{\text{data}}$ with respect to observed transport volume Ξ_h^{data}. Setting $t_h = (\bar{F}_h)^{\sigma_1}$ in line with equation 7 and using fitted values of fuel efficiency $\hat{\bar{F}}_h^{\text{data}}$ in equation 35, we get the following pass-through equation

$$\ln t_h \;=\; \sigma_0 \;+\; \sigma_1^{\text{data}} \cdot \ln \hat{\bar{F}}_h^{\text{data}} \;+\; \zeta_h \tag{36}$$

where σ_1^{data} captures the pass-through elasticity between unit fuel consumption and transport cost.

When using model-generated fuel efficiency to estimate scale effects, α_1 describes the relationship between fuel efficiency $\bar{F}_h$ and transport work $\sum_L (\Xi_{L,h} d_{L,h})$. This change is made to provide a direct correspondence between fuel elasticity and the slope of the fuel efficiency curve shown in Figure 6a. Fuel elasticity can simply be estimated via the

following reduced form equation

$$\ln \bar{F}_h^{\text{model}} \;=\; \alpha_0 \;+\; \alpha_1^{\text{model}} \cdot \ln \sum_L (\Xi_{L,h} d_{L,h}) \;+\; \upsilon_h. \tag{37}$$

while pass-through estimates are obtained from

$$\ln t_h \;=\; \sigma_0 \;+\; \sigma_1^{\text{model}} \cdot \ln \hat{\bar{F}}_h^{\text{model}} \;+\; \zeta_h \tag{38}$$

parameter σ_1^{model} in equation 38. Detailed results of direct and decomposed scale effect estimates are presented in Table 9 and summarized in Table 1. Despite the difference across all three approaches, scale effect estimates are remarkably similar. While larger than estimates reported in GWZ, this appears more related to the underlying data than a difference in methodology as seen when comparing columns 1 and 2 of Table 1. The model-based scale effects are therefore in line with the existing estimates in the literature and can be decomposed into a negative fuel elasticity and a positive pass-through elasticity as assumed by the model.

When conducting policy counterfactuals, I compute endogenous scale effects λ_1^{model} using equations 37 and 38. Whenever there is a change in cost fundamentals $\hat{\tau}$, shippers respond by reallocating their vessels according to equations 12 and 14, resulting in an updated unit fuel consumption $\bar{F}_h^{\text{model, new}}$. Scale effects therefore need to be re-estimated to capture the changed policy environment. Plugging $\bar{F}_h^{\text{model, new}}$ into equations 37 and 38, I get an updated scale effect $\lambda_1^{\text{model, new}} = \alpha_1^{\text{model, new}} \cdot \sigma_1^{\text{model, new}}$.

Note that these elasticities can also be used to update transport cost matrix A in case of a change in cost fundamentals $\hat{\tau}$. A carbon tax that raises unit fuel consumption on EU legs by 10% would for example be expected to raise transport cost by $10 \cdot \sigma_1^{\text{model}}$ percent.

Table 9: Scale effects in the maritime shipping - detailed

Approach 1: Replicate Ganapati et al. (2024)

Dependent variable:	transport volume	transport cost
Stage:	IV	Scale elasticity
Constant	5.819***	3.362
	(2.190)	(3.059)
IV	0.4229***	
	(0.1512)	
log dist	-0.1654***	0.8576***
	(0.0588)	(0.0521)
log transport volume (λ_1)		-0.4328*
		(0.2577)
R^2	0.007	0.595
Observations	1,764	1,764

Approach 2: Decomposition of scale effects

	observed fuel efficiency		**model-implied fuel efficiency**	
Dependent variable:	fuel efficiency	transport cost	fuel efficiency	transport cost
Stage:	Fuel elasticity	Pass-through	Fuel elasticity	Pass-through
Constant	4.456***	4.139***	6.892***	2.090***
	(0.1522)	(0.1686)	(0.2161)	(0.2029)
log transport volume (α_1^{data})	-0.3223***			
	(0.0141)			
log fuel efficiency (σ_1)		1.728***		1.782***
		(0.1607)		(0.0929)
log transport work (α_1^{model})			-0.2396***	
			(0.0107)	
R^2	-1.366	0.062	0.220	0.173
Observations	1,764	1,764	1,764	1,764

The table shows scale effect estimates across different approaches. Direct estimates in approach 1 follow Ganapati et al. (2024) and rely on equations 34. Fuel elasticity estimates are based on equations 35 and 37, while pass-through elasticity uses equations 36 (column 2) and 38 (column 4). Scale elasticity of approach 1 is equal to coefficient λ_1. Scale effects from a decomposition into fuel and pass-through elasticities can be obtained as $\lambda_1^{\text{data}} = \alpha_1^{\text{data}} \cdot \sigma_1$ and $\lambda_1^{\text{model}} = \alpha_1^{\text{model}} \cdot \sigma_1$ for observed and model-generated fuel efficiency respectively. Note that column 1 in approach 2 reports the second-stage estimates of a 2SLS regression in which transport volume is instrumented by the same geography-based IV as used in approach 1. Further details are available in appendix D.4. Significance codes: ***: 0.01, **: 0.05, *: 0.1.

D.5 Pass-through estimation using external data

This section presents pass-through estimates using external data. Bunker fuel prices are obtained from S&P Global Commodity Insights and cover quarterly prices of 18 bunker types listed on different exchanges around the globe. Quarterly crude oil prices are obtained from the FED St Louis online repository (available at: https://fred.stlouisfed.org/series/POILBREUSDM#0). Weekly container freight rates for a select number of shipping routes are collected from Drewry and aggregated to the quarter level. Combining the data creates a panel of freight rates and fuel cost at the route-quarter-bunker (r, t, p) type level.

To estimate pass-through elasticities, I run the following simple regression

$$\ln(\text{freight rate})_{r,t} \;=\; \sigma_1^{\text{ext}} \cdot \ln(\text{price bunker})_{p,t} \;+\; FE_p \;+\; \varepsilon_{r,t,p}$$

where σ_1^{ext} captures the pass-through elasticity between bunker fuel and freight rates. Bunker type fixed effects FE_p control for unobserved differences across bunker exchanges and bunker types, such that σ_1^{ext} is only identified from variation within bunker types over time. To mitigate endogeneity concerns related to unobserved cost shifters which both affect freight rates and bunker prices, I use the global crude oil price as an instrument for bunker prices and estimate pass-through elasticities using 2SLS. Bunker fuel is a direct derivative of crude oil. Bunker prices therefore strongly respond to price changes in crude oil. As ships cannot directly use crude oil to run their engines, it can only affect freight rates via bunker prices, making it a suitable instrument.

Regression results for the World Container Index and routes between Shanghai-LA, LA-Shanghai and Shanghai-Rotterdam are reported in Table 10. Estimated pass-through elasticities range from 0.78 to 1.37 which is slightly below but still within range of the model-implied pass-through rates reported in Table 9.

Table 10: Pass-through estimates using external data

Dependent Variable:	log freight rate			
Route:	WCI	SH to LA	LA to SH	SH to RT
Variables				
log bunker price	1.115***	0.9805***	0.7823***	1.371***
	(0.2955)	(0.3388)	(0.1401)	(0.3298)
Instrument				
log crude price	yes	yes	yes	yes
Fixed-effects				
bunker type	Yes	Yes	Yes	Yes
Fit statistics				
Observations	607	607	607	607
R^2	0.28965	0.22464	0.44535	0.29806
Within R^2	0.27565	0.20913	0.4425	0.28422

The table presents second-stage estimates of pass-through rates obtained estimated via 2SLS. In all regressions, crude oil prices serve as instruments for bunker prices. WCI refers to Drewry World Container Index. SH, LA and RT stand for Shanghai, Los Angeles and Rotterdam respectively. Standard errors are clustered at the year-quarter level. Significance codes: ***: 0.01, **: 0.05, *: 0.1.

E Additional results

E.1 Counterfactual procedure

The algorithm presented below describes the iterative procedure that is used to find a new equilibrium after a policy-induced change in transport cost $\hat{t_{kl}}$ in the network.

Table 11: Algorithm for policy counterfactuals

Steps:	Model equation
function POLICY COUNTERFACTUAL $(X_0, \Xi_0, \hat{t})$	→ Find new equilibrium
Initialize baseline trade X_0, traffic Ξ_0 and transport cost A_0	
Compute baseline allocation of vessels $\bar{F}_0 = \bar{F}(\Xi_0)$	→ Equations 12 and 14
Compute baseline scale elasticity $\lambda_0 = \lambda(\bar{F}_0)$	→ Equations 37 and 38
Initialize change in fundamentals $\hat{\tau}$	→ Example: EU carbon tax
Update cost matrix $A_1 = A(\Xi_0; \lambda_0; \hat{\tau})$	→ See appendix D.4
Compute $B_1 = (I - A_1)^{-1}$	
Initialize loop: diff = inf, tol = ϵ	
while $diff < tol$ **do**	
Update trade flows $X_1 = X(B_1)$	→ Solve equations 25 to 29
Update traffic $\Xi_1 = \Xi(X_1, A_1, B_1)$	→ Equation 6
Update allocation $\bar{F}_1 = \bar{F}(\Xi_1)$	→ Equations 12 and 14
Update scale elasticity $\lambda_1 = \lambda(\bar{F}_1)$	→ Equations 37 and 38
Update cost matrix $A_2 = A(\Xi_1; \lambda_1)$	→ See appendix D.4
Compute $B_2 = (I - A_2)^{-1}$	
Compute $diff = \sum_{ij}(B_2 - B_1)^2$	
Update $A_1 = A_2$ and $B_1 = B_2$	
Return final X_1, Ξ_1, $\bar{F}_1$	
Compute change in real wages $\hat{\mathrm{w}}_i / \hat{P}_i$	→ See appendix E.1
Compute change in transport emissions	→ Follow equation 17

E.2 Benchmark estimates - heterogeneity across countries

Figure 15: Impact of 10% EU carbon tax across countries

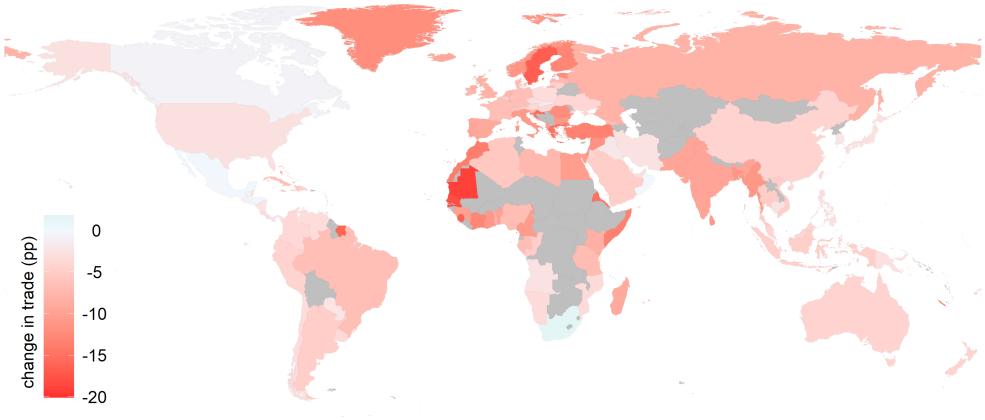

(a) Containerized trade (change in pp)

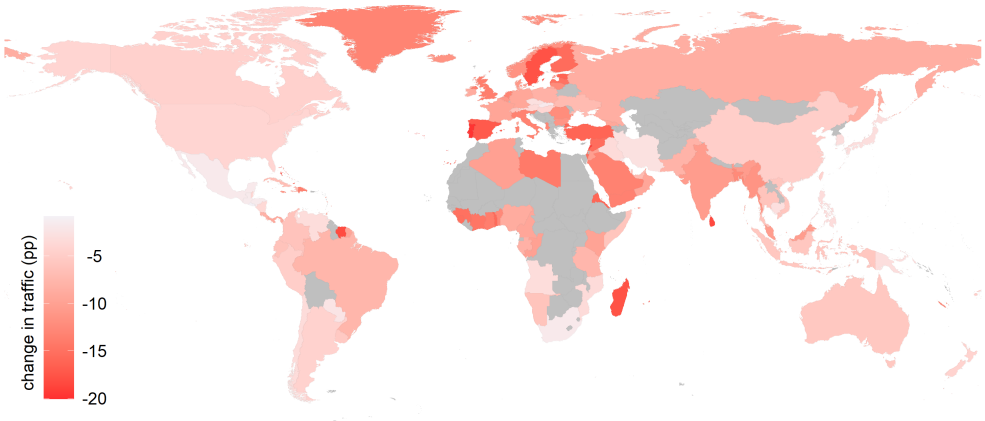

(b) Containerized traffic (change in pp)

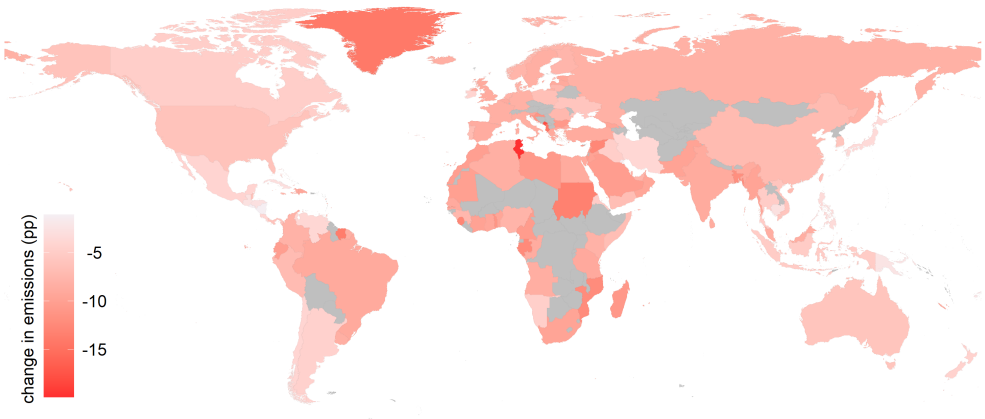

(c) Transport emissions (change in pp)

E.3 Alternative fleet composition

Table 12: Impact of 10% EU carbon tax under alternative fleet compositions

	35% clean fleet		45% clean fleet		100% clean fleet	
idle rate	$\eta = 0$	$\eta = 0.2$	$\eta = 0$	$\eta = 0.2$	$\eta = 0$	$\eta = 0.2$
change traffic (pp)	-9.62	-9.48	-9.60	-9.46	-9.90	-9.99
change efficiency (pp)	17.00	4.95	17.60	4.50	17.62	15.71
change emissions (pp)	2.76	-7.82	3.22	-8.17	2.82	0.93

The table shows counterfactual estimates for a 10% EU carbon tax under alternative fleet compositions. Columns 1 and 2 show results of the benchmark setting for a fully inelastic ($\eta = 0$) and elastic ($\eta = 0.2$) fleet. Column pairs 3-4 repeats the analysis in a setting where the share of clean vessels within the fleet is raised by 10pp. Column pair 5-6 shows results for a fleet that only consists of clean ships.

E.4 Allocation distortion from fuel standards

In this section, I present a counterfactual exercise which is meant to simulate EU fuel standards without the need to observe fuel consumption data across different fuel types in 2018. Instead of trying to capture an optimal fuel choice for each clean vessel in the EU, I simplify the setting by assuming that every clean vessel is required to adopt a new fuel type to meet the new standard.

The new fuel type is assumed to have a lower emission factor and a higher cost than the current one. Adopting it therefore decreases the emission intensity which helps to comply with EU standards but also increases the cost of clean vessels when operating on EU segments, making it less attractive for shippers to allocate clean capacity to EU routes. To see this, I adopt the notation of the theoretical framework. The policy intervention can be represented by an EU-specific cost shock that raises the baseline cost parameter of clean vessels from δ_c to $\delta_c^{EU} = \delta_c * s_c^{EU}$ where s_c^{EU} represents a fuel surcharge compared to the baseline setting. As $\delta_c^{EU} > \delta_c^{nonEU}$, routes featuring EU segments achieve a lower ranking in the allocation of clean capacity in equation 12 because the potential savings on EU segments have decreased ($|\delta_c^{EU} - \delta_d| < |\delta_c^{nonEU} - \delta_d|$). On the margin, this incentivizes shippers to allocate clean capacity to non-EU routes instead, especially if fuel surcharge s_c^{EU} becomes large.

To comply with the policy, shippers have to ensure that the fuel mix M^{EU} defined as

the share of total EU fuel consumption coming from clean vessels does not fall below the baseline level M_0^{EU}. To do so, shippers can distort the allocation of clean capacity in favour of EU routes by keeping the share of clean capacity identical to the baseline.

In this setting, an allocation distortion occurs whenever the EU fuel mix falls below its baseline level. To start with, I investigate whether this is the case when combining fuel standards with an EU carbon tax of 1%. Panel 1 of Table 13 presents fuel mix estimates for varying levels of fuel surcharge s_c^{EU}. Results indicate that the fuel mix generally falls with surcharge level s_c^{EU} but only substantially deviates from its baseline level at surcharge rates of 50% and above. Below that, counterfactual allocations of clean capacity in the EU are quasi-compliant with the fuel standard. The maximum distortion at a 10% surcharge level for example amounts to merely 4% when comparing the unconstrained allocation in panel 1 to the constrained allocation in panel 2 which is fully compliant with the fuel standard.

As EU fuel standards in practice are not expected to raise the cost of EU vessels by 50% above their baseline level, I assume that they do not play a decisive role for the environmental gains of carbon policy.

Table 13: Simulated allocation distortion in the EU

fuel surcharge	0%	1%	10%	50%	95%
no allocation barrier:					
fuel mix	0.17	0.17	0.17	0.15	0.04
route share	0.09	0.09	0.09	0.06	0.01
fixed EU route share:					
fuel mix	0.17	0.17	0.18	0.20	0.24
route share	0.09	0.09	0.09	0.09	0.09
allocation distortion:					
distortion fuel mix	1.00	1.00	1.04	1.38	6.29
distortion route share	1.00	1.01	1.04	1.37	9.26

The table compares counterfactual outcomes of a cost shock emanating from a fuel surcharge on shipping routes that pass through the EU as explained in section 6. Fuel mix and route share refer to the share of total EU fuel consumption and total number of EU routes that are served by clean vessels. Panel 1 shows the unrestricted response of shippers to the shock while panel 2 shows the response when shippers are forced to serve EU routes with the same share of clean vessels as in the baseline. Panel 3 compares them by dividing restricted and unrestricted outcomes.

E.5 Production emission - details

Figure 16: Impact of 10% EU carbon tax on production emissions

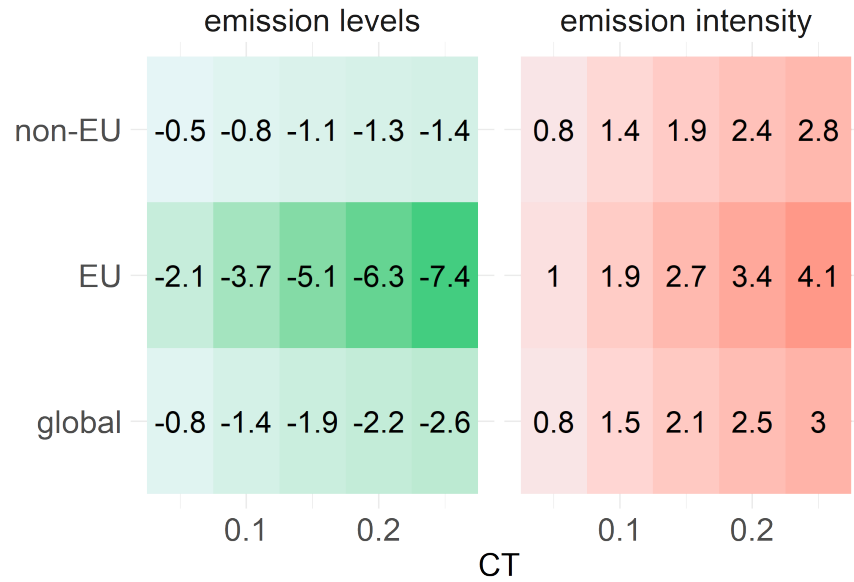

(a) production emissions of traded goods (change in pp)

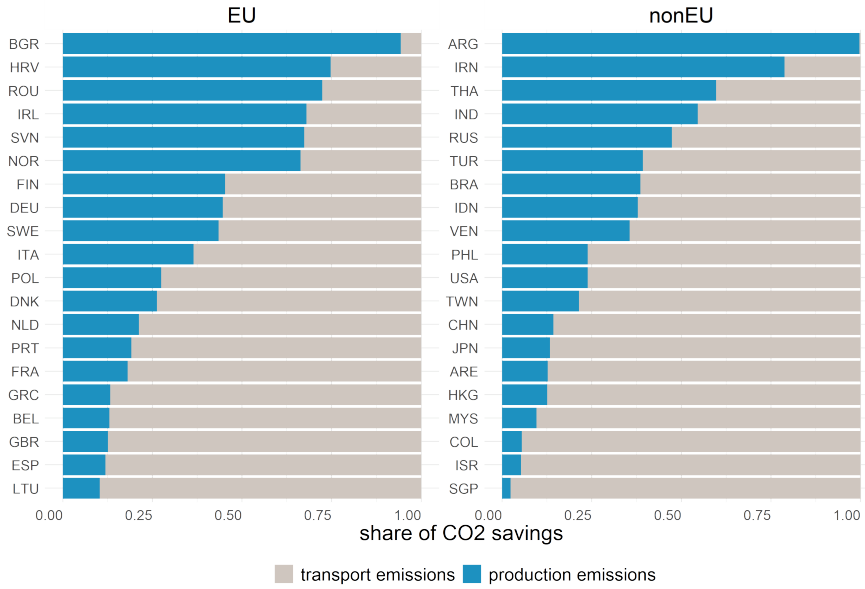

(b) share of production emissions in total CO2 savings

Note: This Figure shows the change in production emissions in response to a EU carbon tax on maritime shipping. Panel a) summarizes the percentage point change in emission levels and intensity for various tax scenarios and regions. A 10% EU carbon tax for example reduces global emission levels by 1.4% but raises global emission intensity (CO2 per USD) by 1.5%. Panel b) instead compares the relative share of emission savings from production and transport emissions across countries.

E.6 Welfare results

Figure 17: Welfare ratio across countries

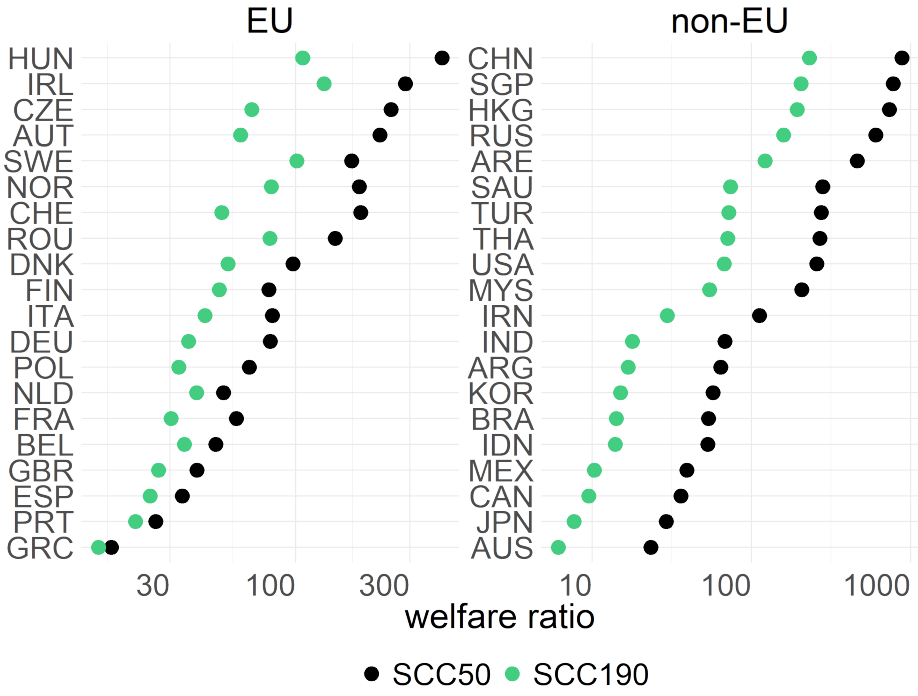

Note: This Figure compares the change in environmental and non-environmental welfare components, summarized as a 'welfare ratio', across countries following a 10% EU carbon tax. Welfare ratios are defined as the change in real income divided by the sum of social carbon benefits (in USD at given social cost of carbon (SCC) and tax revenue (in EU countries).

Table 14: Welfare estimates across countries

				unweighted SCC		weighted SCC	
	country	Δ real wage	Δ CO2	Δ welfare	ratio	Δ welfare	ratio
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1	AGO	-206.71	-15.34	-198.33	-24.67	-205.32	-148.33
2	ALB	-107.59	-6.80	-99.21	-12.84	-107.52	-1553.78
3	ARE	-3326.88	-373.65	-3318.50	-397.04	-3319.68	-462.12
4	ARG	-587.34	-30.30	-578.96	-70.10	-578.18	-64.08
5	AUS	-798.11	-68.15	-789.73	-95.25	-763.91	-23.34
6	AUT	-2389.60	-28.85	-2381.22	-285.18	-2378.95	-224.44
7	BEL	-6975.63	-694.38	-6836.18	-50.02	-6825.74	-46.54
8	BEN	-34.05	-13.71	-25.68	-4.06	-33.94	-292.18
9	BGD	-661.00	-42.17	-652.62	-78.89	-658.25	-241.00
10	BGR	-897.80	-52.54	-888.90	-100.93	-895.51	-392.57
11	BHR	-462.55	-43.02	-454.17	-55.20	-462.55	-55.20
12	BHS	-58.86	-67.89	-50.48	-7.02	-58.67	-312.86
13	BLZ	-5.38	-0.34	3.00	-0.64	-5.33	-107.57
14	BMU	-9.92	-0.13	-1.54	-1.18	-9.82	-100.56
15	BRA	-1721.66	-183.97	-1713.28	-205.47	-1689.65	-53.79
16	BRB	-34.94	-2.39	-26.56	-4.17	-34.85	-376.02
17	BRN	-30.94	-0.27	-22.56	-3.69	-30.52	-74.89
18	CAN	-1383.25	-112.13	-1374.87	-165.08	-1344.96	-36.13
19	CHE	-3111.35	-17.02	-3102.97	-371.32	-3094.69	-186.76
20	CHL	-439.67	-19.50	-431.29	-52.47	-433.83	-75.21
21	CHN	-23206.80	-1415.50	-23198.42	-2769.59	-23180.42	-879.77
22	CIV	-257.04	-30.91	-248.66	-30.68	-256.57	-548.42

Table 14: Welfare estimates across countries

				unweighted SCC		weighted SCC	
	country	Δ real wage	Δ CO2	Δ welfare	ratio	Δ welfare	ratio
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
23	CMR	-133.44	-11.54	-125.06	-15.93	-132.99	-295.13
24	COG	-105.15	-10.72	-96.77	-12.55	-104.96	-548.23
25	COL	-232.36	-132.06	-223.98	-27.73	-227.03	-43.62
26	CRI	-106.56	-15.54	-98.18	-12.72	-105.79	-137.81
27	CUB	-140.05	-16.05	-131.67	-16.71	-139.01	-134.81
28	CYM	-24.79	-1.00	-16.41	-2.96	-24.64	-170.62
29	CYP	-248.49	-65.61	-230.54	-13.84	-238.55	-25.01
30	CZE	-1492.23	-25.08	-1483.85	-178.09	-1486.26	-249.78
31	DEU	-17721.52	-1011.55	-17582.52	-127.49	-17495.16	-78.29
32	DNK	-2772.20	-112.47	-2744.11	-98.69	-2743.73	-97.37
33	DOM	-178.90	-94.49	-170.52	-21.35	-177.75	-154.79
34	DZA	-862.34	-17.94	-853.96	-102.91	-859.69	-325.63
35	ECU	-70.12	-28.95	-61.74	-8.37	-68.50	-43.34
36	EGY	-633.99	-1276.27	-625.61	-75.66	-629.68	-147.19
37	ERI	-8.13	-0.15	0.25	-0.97	-8.01	-67.90
38	ESP	-8123.17	-1084.84	-7904.75	-37.19	-7881.99	-33.68
39	EST	-644.09	-77.83	-628.11	-40.31	-635.77	-77.37
40	FIN	-2268.23	-166.86	-2236.87	-72.34	-2238.87	-77.25
41	FJI	-19.26	-1.57	-10.88	-2.30	-19.19	-282.96
42	FRA	-11333.64	-741.99	-11187.59	-77.60	-11133.53	-56.64
43	GAB	-96.65	0.13	-88.27	-11.53	-95.97	-143.96
44	GBR	-15494.33	-1789.46	-15149.84	-44.98	-15094.58	-38.76
45	GEO	-68.21	-5.56	-59.83	-8.14	-68.14	-1036.75
46	GHA	-396.92	-21.95	-388.54	-47.37	-396.24	-585.19
47	GIN	-100.95	-4.97	-92.57	-12.05	-100.85	-988.22
48	GMB	-9.28	-1.95	-0.90	-1.11	-9.23	-209.90
49	GRC	-1595.00	-487.14	-1497.64	-16.38	-1501.63	-17.08
50	GRL	-11.25	-1.05	-2.87	-1.34	-11.19	-200.39
51	GTM	-50.99	-11.62	-42.61	-6.09	-50.17	-62.42
52	HKG	-6160.37	-300.33	-6152.00	-735.20	-6160.37	-735.20
53	HND	-31.87	-3.82	-23.49	-3.80	-31.38	-64.36
54	HRV	-909.62	-49.39	-899.21	-87.39	-906.31	-275.10
55	HTI	-22.16	-2.12	-13.78	-2.64	-22.05	-195.34
56	HUN	-1596.75	-26.89	-1588.37	-190.56	-1592.83	-407.42
57	IDN	-993.22	-65.13	-984.84	-118.53	-974.52	-53.13
58	IND	-6056.84	-514.51	-6048.46	-722.85	-5967.93	-68.12
59	IRL	-5943.95	-78.92	-5923.61	-292.21	-5923.24	-287.04
60	IRN	-711.88	-62.94	-703.50	-84.96	-705.55	-112.58
61	IRQ	-600.23	-37.13	-591.85	-71.63	-596.12	-146.28
62	ISL	-184.32	-27.33	-172.66	-15.80	-180.55	-48.85
63	ISR	-1571.56	-136.82	-1563.18	-187.56	-1565.01	-240.01
64	ITA	-11343.24	-672.50	-11239.29	-109.12	-11201.62	-80.09
65	JAM	-58.18	-51.61	-49.80	-6.94	-57.89	-199.59
66	JOR	-257.29	-10.60	-248.91	-30.71	-256.71	-445.14
67	JPN	-3351.55	-226.29	-3343.17	-399.99	-3236.78	-29.20
68	KEN	-92.85	-10.86	-84.47	-11.08	-92.11	-125.23
69	KHM	-75.59	-2.34	-67.21	-9.02	-75.29	-249.61
70	KOR	-2662.43	-384.16	-2654.05	-317.74	-2615.99	-57.33
71	KWT	-913.21	-14.44	-904.83	-108.99	-911.59	-563.53
72	LBN	-414.81	-152.29	-406.43	-49.50	-414.21	-690.97
73	LBY	-133.54	2.17	-125.16	-15.94	-133.20	-391.07
74	LKA	-290.75	-405.28	-282.37	-34.70	-289.27	-197.19
75	LTU	-406.73	-64.61	-383.06	-17.18	-390.26	-24.70
76	LVA	-377.88	-40.14	-362.14	-24.02	-369.79	-46.73
77	MAC	-104.69	-2.02	-96.31	-12.49	-104.54	-664.31
78	MAR	-1319.84	-663.52	-1311.47	-157.52	-1317.90	-677.58
79	MDG	-73.24	-12.66	-64.86	-8.74	-72.87	-198.82
80	MEX	-906.15	-24.62	-897.77	-108.14	-883.07	-39.26
81	MLT	-1188.26	-248.16	-1126.86	-19.35	-1134.83	-22.24
82	MMR	-41.20	-0.99	-32.82	-4.92	-39.47	-23.75
83	MNE	-72.87	-7.67	-64.50	-8.70	-72.85	-2797.94
84	MOZ	-47.89	-7.54	-39.52	-5.72	-47.51	-125.89
85	MRT	-89.72	-4.06	-81.34	-10.71	-89.65	-1279.19
86	MUS	-127.97	-53.99	-119.60	-15.27	-127.65	-388.65
87	MYS	-1553.09	-847.57	-1544.71	-185.35	-1545.59	-206.92
88	NAM	-45.73	-4.55	-37.35	-5.46	-45.35	-120.78
89	NCL	-62.94	-6.18	-54.56	-7.51	-62.67	-237.54

Table 14: Welfare estimates across countries

				unweighted SCC		weighted SCC	
	country	Δ real wage	Δ CO2	Δ welfare	ratio	Δ welfare	ratio
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
90	NGA	-457.50	-16.31	-449.12	-54.60	-449.76	-59.11
91	NIC	-31.86	-2.06	-23.48	-3.80	-31.65	-152.37
92	NLD	-11506.78	-1093.29	-11292.83	-53.78	-11277.07	-50.09
93	NOR	-3302.64	-111.93	-3284.89	-186.05	-3284.75	-184.54
94	NZL	-464.00	-22.38	-455.62	-55.38	-459.00	-92.84
95	OMN	-248.15	-147.07	-239.77	-29.62	-247.03	-221.35
96	PAK	-496.33	-104.75	-487.95	-59.23	-491.25	-97.56
97	PAN	-123.00	-171.30	-114.62	-14.68	-122.02	-126.55
98	PER	-298.61	-32.23	-290.23	-35.64	-295.11	-85.30
99	PHL	-416.59	-19.41	-408.21	-49.72	-412.68	-106.68
100	PNG	-52.00	-0.79	-43.62	-6.21	-51.65	-148.89
101	POL	-2669.09	-124.83	-2633.46	-74.91	-2627.44	-64.08
102	PRT	-1674.97	-262.09	-1607.31	-24.76	-1610.82	-26.11
103	PRY	-40.21	-0.09	-31.83	-4.80	-39.61	-67.01
104	PYF	-30.25	-2.17	-21.87	-3.61	-30.06	-157.61
105	QAT	-694.69	-15.59	-686.31	-82.91	-692.67	-344.59
106	ROU	-2177.25	-140.99	-2158.60	-116.75	-2162.36	-146.20
107	ROW	-5766.71	-412.62	-5758.33	-688.22	-5739.74	-213.86
108	RUS	-5057.75	-287.14	-5049.37	-603.61	-5057.75	-603.61
109	SAU	-2541.26	-509.24	-2532.88	-303.28	-2532.19	-280.21
110	SDN	-54.88	-17.22	-46.50	-6.55	-52.97	-28.76
111	SEN	-150.41	-42.24	-142.04	-17.95	-150.14	-552.14
112	SGP	-6523.59	-1143.89	-6515.21	-778.55	-6523.59	-778.55
113	SLE	-49.14	-2.58	-40.76	-5.86	-48.97	-302.39
114	SLV	-20.29	-1.12	-11.91	-2.42	-19.65	-31.80
115	SOM	-0.14	-1.22	8.24	-0.02	-0.12	-8.23
116	SRB	-511.38	-81.51	-503.00	-61.03	-511.15	-2197.53
117	SUR	-44.91	-4.42	-36.53	-5.36	-44.80	-396.54
118	SVK	-766.57	-13.21	-758.20	-91.49	-763.41	-242.04
119	SVN	-579.59	-28.76	-569.36	-56.66	-576.43	-183.04
120	SWE	-7903.22	-270.66	-7860.00	-182.87	-7856.98	-170.95
121	SYC	-33.74	-5.07	-25.36	-4.03	-33.68	-499.31
122	SYR	-157.13	-6.22	-148.75	-18.75	-156.02	-142.71
123	TGO	-37.09	-30.22	-28.71	-4.43	-36.93	-238.65
124	THA	-3348.66	-194.30	-3340.29	-399.64	-3336.22	-269.13
125	TTO	-89.43	-3.06	-81.05	-10.67	-88.52	-97.78
126	TUN	-1038.69	-71.32	-1030.31	-123.96	-1038.00	-1512.56
127	TUR	-3926.36	-531.63	-3917.98	-468.59	-3912.03	-274.08
128	TWN	-991.24	-299.71	-982.86	-118.30	-991.24	-118.30
129	TZA	-111.41	-7.80	-103.03	-13.30	-110.77	-174.79
130	UKR	-694.16	-96.11	-685.78	-82.84	-693.30	-814.46
131	URY	-73.43	-2.12	-65.05	-8.76	-73.43	-8.76
132	USA	-15435.45	-764.84	-15427.07	-1842.13	-15375.34	-256.81
133	VEN	-218.67	-3.39	-210.29	-26.10	-212.60	-35.98
134	VNM	-819.79	-223.08	-811.41	-97.84	-816.90	-283.50
135	WSM	-4.92	-0.18	3.46	-0.59	-4.89	-186.32
136	YEM	-40.61	0.64	-32.23	-4.85	-39.88	-56.30
137	ZAF	71.37	19.51	79.74	8.52	77.96	10.83

This table shows welfare estimates from a 10% EU carbon tax and a social cost of carbon (SCC) of 50 USD. Columns 2 shows the change in real wages in million USD while column 3 depicts the change in carbon emissions in thousand tons of CO2. Emission savings are distributed equally across countries in columns 4 and 5, while columns 6 and 7 weight the damage across countries following the approach of Shapiro (2016). Welfare changes in columns 4 and 6 are measured in million USD. Welfare ratios in columns 5 and 7 divide the changes in real wages by the gains from emission savings. Welfare changes and ratios account for tax revenues associated with levying carbon taxes for traffic in EU waters for respective countries. Tax revenues are hereby treated as lump-sum payments to consumers.