The Behavioral Impact of Sulfur Fuel Content Regulations in Maritime Shipping

William Troske*

November 1, 2026

Link to Current Version

Abstract

In 2012, the US and Canada implemented a regulation called the ECA (Emission Control Area) in their territorial and economic waters to reduce emissions from ocean-going ships. The ECA required ships to use low-sulfur fuel, increasing a ship's fuel costs. In transportation, an increase in fuel costs can lead to changes in shipping supply characteristics such as operating speed, ship size, and trip frequency. This paper aims to quantify the extent of adjustment in maritime shipping supply characteristics in response to the ECA regulation. I first present a stylized model to highlight how the ECA affects shipping supply characteristics by changing the cost of fuel. I then test these predictions using port entrance and exit data from the US Army Corps of Engineers. I find that for the aggregate cost increase, treated routes saw ships that were, on average, about 15 percent larger in gross tons and had 25 more trips per quarter compared with similar, untreated routes, and ships on treated routes moved 16 percent slower than predicted in the imputed counterfactual. The size of the cost change matters, however, for the response to ship size.

Keywords: emission control areas, maritime shipping, fuel regulation, slow steaming

JEL Codes: Q53, Q58, R41, L91

*Department of Agricultural and Resource Economics, University of California, Davis. Email: wtroske@ucdavis.edu

1 Introduction

Maritime shipping contributes to an estimated 13% of global SO_x and NO_x pollution (International Maritime Organization, 2014) and 10% of US SO_x and NO_x pollution (US EPA, 2009). At the same time, shipping accounts for 70% of global trade volume (Kalouptsidi, 2014) and 71% of US international trade volume (International Trade Administration, 2024).

Concerned about high pollution levels, Canada and the US created the North American ECA (Emission Control Area) to restrict sulfur emissions, but at a potentially high cost to ships. The ECA requires ships operating within US and Canadian waters (up to 200 nautical miles from the coast) to use low-sulfur fuel to reduce SO_x emissions. The EPA estimated the ECA would prevent 3,700 to 8,300 premature deaths while increasing fuel costs by $145 per tonne (about 30%) over conventional marine fuels (US EPA, 2009). While the social benefits are large, the added fuel costs are a significant burden for operators, since fuel represents 47% of a voyage's total costs (Stopford, 2009).

Such costly regulations in a large industry may have lasting effects on the supply of transportation. In the maritime industry, companies deliver goods in the quantities and frequencies demanded while minimizing costs (Stopford, 2009). An increase in fuel costs can lead companies to either slow down or use larger ships. Slowing down uses less fuel for any ship, and larger ships are more fuel efficient per unit transported (Stopford, 2009); however, slowing down or using larger ships can disrupt the service frequency along a set shipping route (Notteboom and Vernimmen, 2009). Taken together, these margins of adjustment allow companies to absorb the cost of regulation rather than pass the costs to consumers.

This paper quantifies the impacts of the ECA regulation on maritime shipping characteristics (ship size, trip frequency, and ship speed). To develop predictions, I first present a stylized model of ships operating on a route showing how the ECA affects shipping supply characteristics by changing fuel costs. I then test these predictions using port entrance and exit data from the US Army Corps of Engineers. To estimate the ECA's effect on ship size and frequency, I develop a novel route-level identification strategy that compares treated and untreated routes using a difference-in-differences approach. Because speed is not observed for the control group, I estimate the ECA's effect on ship speed using an imputation approach following Borusyak et al. (2024). After estimating the aggregate effects, I examine how the ECA's impact varies with route length within the ECA, as route length proxies for heterogeneous cost differences.

To demonstrate how an increase in fuel costs can change shipping supply characteristics, I construct a stylized model of shipping drawing on Cullinane and Khanna (1999) and Peters (2024). A company operating a ship decides whether to serve the route, and then the operating speed and the number of trips. Since larger ships are more fuel efficient, a fuel price increase raises the marginal cost for all ships, but raises it the most for the smallest ships. Companies respond by slowing ships down, which requires additional trips to maintain service frequency. For small cost increases, all ships can

slow down enough to offset the cost increase, and the average ship size on a route does not change. For a moderate cost increase, slower speeds are not enough for the smallest ships to offset the cost increase, so they exit the route, leading to a larger average ship size on the route. For the largest shocks, slower speeds do not fully offset the cost increase. The resulting drop in quantity demanded leaves the largest ships underutilized, prompting them to leave and lowering the average ship size on the route.

My empirical results are consistent with the stylized model's predictions. I find that for the aggregate cost increase, treated routes saw ships that were, on average, about 15% larger in gross tons and had 25 more trips per quarter compared to untreated routes, and the ships on treated routes traveled 16% slower than predicted in the imputed counterfactual. In terms of heterogeneous treatment effects, I find that routes with shorter sections in the ECA see little change in average ship size, routes with moderate sections in the ECA see an increase in average ship size, and routes with longer sections in the ECA see a decrease in ship size, matching the predicted effect of the model.

This research contributes to the literature in several ways. First, by providing the first reduced-form estimates of how the ECA led to ships slowing down, companies switching to larger ships, and companies increasing frequency, I bridge research from two fields. One set of the research in economics focuses on short-run changes in ship speed within the ECA (Klotz and Berazneva, 2022, 2025; Moore et al., 2018; Adland et al., 2017)[1]. This strand uses reduced-form techniques to estimate the impact on ship speed in the first months after the ECA. The other set of research comes from the transportation literature and examines companies' choices about ship speed, frequency, ship size, and ports. These modeling papers use dynamic optimization simulations to understand how companies trade off different characteristics when faced with either higher fuel prices in general or the ECA. For example, Notteboom and Vernimmen (2009) study how higher fuel prices affect speed, size, and frequency decisions between China and Europe. Lan et al. (2023) and Zavitsas et al. (2018) perform simulations of optimal responses to the ECA for a specific route between the US and China. My research complements these modeling studies by using observational data to estimate the supply responses that are built into these simulation models.

Second, this research offers a potential mechanism behind the broader finding that fuel-cost pass-through is small, and trade is inelastic to fuel costs. Freight rates have been found to be relatively inelastic to oil prices, with estimates ranging from 0.1 to 0.3 (Hummels, 2007; Mirza and Zitouna, 2009; Beverelli et al., 2010). At the same time, trade is responsive to transportation costs, with estimates in the range of -3 (Asturias, 2020; Brancaccio et al., 2020; Hummels, 2007; Ganapati et al., 2024; Wong, 2022). Recent work by Brancaccio et al. (2023) finds low pass-through of fuel costs to transportation prices (0.17) and, using the model of Brancaccio et al. (2020), finds that the elasticity of trade to fuel prices is 0.35. My empirical results provide a potential explanation for these findings. My results show that companies offset fuel cost increases by varying ship operating speeds and, on average, using larger ships.

[1] Additional reduced-form research has studied the pollution impact of the ECA (Hansen-Lewis and Marcus, 2025; Zhu and Wang, 2021).

Third, this research shows how regulating one input in a multi-input production function has heterogeneous impacts on producers. For example, I show that the amount of the fuel cost increase matters for which ships stay on a route and which leave. Previous work shows how environmental regulations can have heterogeneous impacts. Work studying the cap-and-trade markets for NO_x and SO_2 shows how different incentives can lead to emissions moving to facilities more exposed to market prices (Fowlie, 2010; Toyama, 2024). In the cement industry, a minimum standard for kiln efficiency, a key technology in cement production, led to higher barriers to entry and more market power for existing cement firms (Ryan, 2012; Fowlie et al., 2016). For the shipping industry, Peters (2024) shows how potential regulations on engine efficiency, emissions taxes, and speeds could lower emissions but would lead ship operators to extend the life of their existing ships, limiting the adoption of new emissions and efficiency technologies. My stylized model shows how changes in fuel costs improve profitability for larger ships while making profitability worse for smaller ships. In addition, my empirical results comparing routes with different exposure levels to the ECA show that heterogeneity in the routes' exposure to the regulation leads to different impacts. Routes with small cost increases do not change average ship size, routes with moderate cost increases see larger average ships, and routes with high cost increases see smaller average ships.

2 Background

2.1 ECA

The North American ECA (emission control area) regulates the sulfur emissions of commercial ships while they operate in the US and Canadian waters. The US and Canada proposed the ECA to the IMO (International Maritime Organization) in 2009. The IMO, the United Nations organization that manages global shipping agreements, approved it in March 2010. The approved area follows the EEZ of the continental US, Hawaii, and southern Canada, extending 200 nautical miles from the shore of the US and Canada except for areas where the US EEZ is limited by other countries' claims, such as around Miami (Figure 1).

The ECA limits sulfur emissions by requiring ships to use fuel with set limits of sulfur when operating in the ECA. The ECA set the maximum allowable sulfur content in fuel at 1% by weight starting August 1st, 2012, then tightening to 0.1% by weight on January 1st, 2015 (US EPA, 2010). The cost of using fuels that met the regulation was estimated to be an additional $145 per ton (US EPA, 2010).

For the US, the US Coast Guard enforces the ECA regulation. The Coast Guard requires documentation of the fuel changeover with the date and time of the changeover, the level of all fuel tanks (both tanks with standard fuel and low sulfur fuel), and certifications from the ship's fuel supplier (U.S. Coast Guard Office of Commercial Vessel Compliance, 2014).

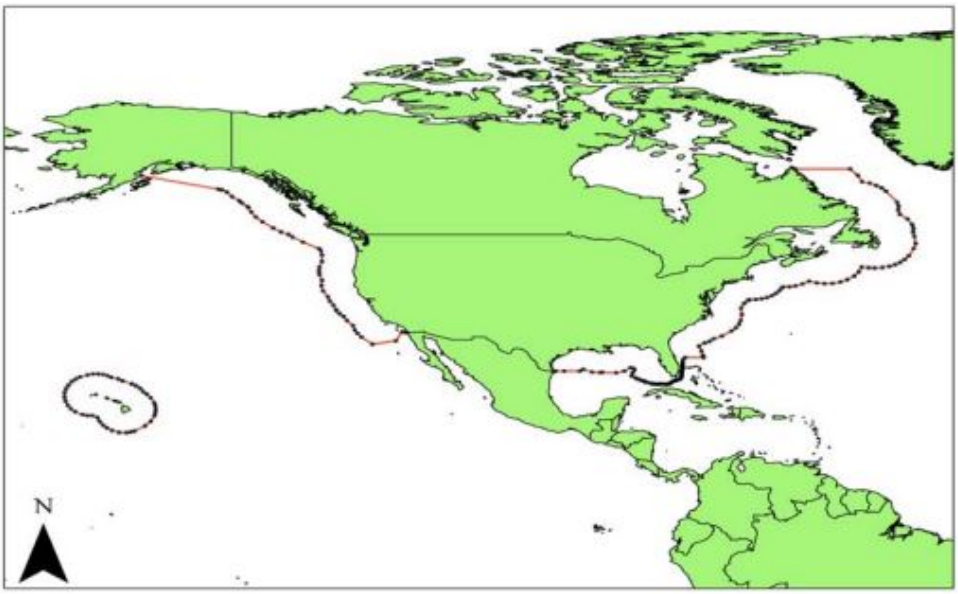


Figure 1: North American ECA region in 2012

Note: Map of North American ECA Boundary (dotted) where ships are required to use low-sulfur fuels. The boundary extends 200 nautical miles from the coast of the US and Canada.

2.2 Container Shipping

To show how the ECA could affect container shipping, I first describe how container shipping operates and present a demonstration model to highlight the potential effects of the ECA.

Containerships operate similarly to buses. A ship travels a set route with pre-planned stops at set times, picking up and dropping off containers at each stop. A ship must be large enough to transport the demanded number of containers and travel fast enough to meet the frequency expectations, often weekly (Notteboom and Vernimmen, 2009).

For companies planning operations on a route, the main variable costs are fuel and capital. Fuel cost represents about 50 percent of the total cost of a single voyage (Stopford, 2009). Companies change fuel costs by operating ships at slower speeds, or they can switch to larger ships. Slower speeds require less fuel, and larger ships are more fuel efficient on a per-unit basis (Stopford, 2009). The cost of capital is the other main cost. Capital represents 42% of the overall cost of the ship. Cullinane and Khanna (1999) find that capital costs also have returns to scale for ship size. Combined, this means larger ships have lower fuel and capital costs per container.

To minimize the costs of fuel and capital, a company chooses ship speed, ship size, and frequency. These parameters are interrelated, and changing one parameter will spill over to affect the optimal choice for all three. Notteboom and Vernimmen (2009) highlights the interrelated nature of these parameters when they study these tradeoffs for the container route between China and Europe. When bunker fuel prices are high, they find that increasing ship speed increases fuel

use in an exponential manner. Slowing ships down to reduce fuel costs, however, slowing down too much means another ship is needed on the route to maintain frequency. Larger ships provide cost savings per container because they are more fuel efficient per container. They operate at higher speeds compared to smaller ships, but port logistics can take more time, and larger ships need enough container demand to make use of the larger ships.

2.3 Sylized Model

To highlight how the ECA acts as a permanent fuel cost shock on ship speed, ship size, and frequency, I build a stylized model of one-period shipping. The model follows the models of Peters (2024) and Cullinane and Khanna (1999). I start with a given ship serving a single route in a quarter, and the ship uses all its slots each trip. The ships produce revenue by transporting containers, which depends on ship size, s (TEU per trip), speed, v (nautical miles /day), and frequency of trips over the quarter, n (trips per quarter). Ships want to go faster to move more containers, but moving faster comes with fuel costs (Equation (1)).

$$\max_{v} \quad \underbrace{pq}_{\text{Revenue}} - \underbrace{nf\alpha s^b dv^2}_{\text{Fuel cost}} - \underbrace{\frac{nrds^c}{v}}_{\text{Size cost}} \quad \text{s.t.} \quad \underbrace{q = \frac{nvs}{d}}_{\text{output constraint}} \tag{1}$$

Revenue comes from shipping q (TEU per quarter) containers at price p ($ per TEU). Fuel costs (Equation (1)) come from the underlying fuel efficiency of the ship αs^b, the fuel usage dv^2, and the price of fuel, f ($ per ton). This relationship comes from the physics of shipping (Stopford, 2009). Fuel efficiency increases with ship size because the volume of a ship, where it stores containers, grows more than the surface area of the hull, which interacts with the ocean and needs fuel to push through. α ($\frac{ton\ days^2}{TEU nm^3}$) represents the efficiency conversion parameter, and s^b, when b is less than 1, represents how ship size increases fuel efficiency for a container. Peters (2024) estimates b to be 0.441. dv^2 represents the fuel usage component. Fuel usage comes from the distance traveled, d (nautical miles). Speed is squared because a ship needs increasing amounts of fuel to maintain higher speeds (Peters, 2024).

The cost of using the ship comes from the literature and captures the opportunity cost of alternatives and the diminishing cost of size (Equation (1)). The price of ship size is r ($ per TEU-day). To capture the diminishing cost of size, ship size scales with c, which is less than 1. Cullinane and Khanna (1999) estimate c to be 0.76. The term d/v converts ship size cost into a cost-per-day parameter.

The production function of shipped quantities captures how speed, frequency, size, and the route interact to produce the quantity of shipped goods. Solving the profit function for optimal speed shows that speed decreases with fuel prices and increases with size, indicating that the result of the ECA will be a slowdown for those ships that continue to serve

the US after the ECA.

$$v^* = \left(\frac{rs^{c-b}}{\alpha f}\right)^{1/3} \tag{2}$$

Through the optimal speed, fuel costs affect the number of ships, the size of the ships, and the overall cost function. Substituting v^* into the production function, with a fixed number of TEUs per quarter, the frequency of trips must increase with rising fuel costs (Equation (3)) since a ship is now moving slower.

$$n = \frac{qd}{sv^*} = \frac{qd}{s}\left(\frac{\alpha f}{rs^{c-b}}\right)^{1/3} \tag{3}$$

Instead of an optimal ship size, I solve for the size at the zero-profit cutoff (Equation (4)). The ships below this cutoff earn negative profit after the increase in fuel costs. Since c and b are less than one and because c is less than twice b, the exponent is negative. When fuel price increases, the zero-profit cutoff for ship size goes up as well.

$$s = \left(\frac{p}{d^2(\alpha f)^{2/3} r^{1/3}\left(1+\frac{1}{r}\right)}\right)^{\frac{3}{c-2b-3}} \qquad \text{(Zero profits)} \tag{4}$$

To solve for the marginal cost as a function of ship size, I plug in the optimal speed. This results in a marginal cost curve that is increasing in ship size, and fuel cost changes the slope of the cost curve (Equation (5)).

$$C(s,q) = 2qd^2 r^{1/3}(\alpha f)^{2/3} s^{\frac{c-2b}{3}-1} \implies \frac{\partial C}{\partial q} = 2d^2 r^{1/3}(\alpha f)^{2/3} s^{\frac{c-2b}{3}-1} \tag{5}$$

Plotting the cost curve using a fixed distribution of ships (Figure 2), all providing a set frequency, the largest ships supply at the lowest cost, and marginal costs increase inversely with ship size. The largest ships carry the most containers and operate at the highest speeds. Smaller ships carry fewer containers and operate at a lower optimal speed; this means even with more ships of that size, those ships supply a lower percentage of the aggregate quantity.

With changing fuel costs due to the ECA, the cost curve does not shift upwards parallel but changes the slope as well (Equation 5), since f makes up the slope of the curve. A fuel cost increase increases costs for all ships (Figure 3). To compensate for the increase in the price of fuel, all ships slow down to lower fuel use. To maintain the frequency, the number of trips must increase. For a small price increase, all ships adjust speed to lower fuel costs (f = \$700 in Figure 3), because prices are slightly higher and most ships are not the marginal ship; they earn higher profits than before the price increase. The average ship size doesn't change much, but frequency increases to compensate for slower ships. For moderate fuel price increases, ships must slow down more to compensate for the price increase (f=\$900 in Figure 3). The speed and quantity decrease enough that the marginal ships no longer make positive profits. The small marginal ships likely leave the route, shifting the average ship size to be larger. The shipping price increase outweighs the quan-

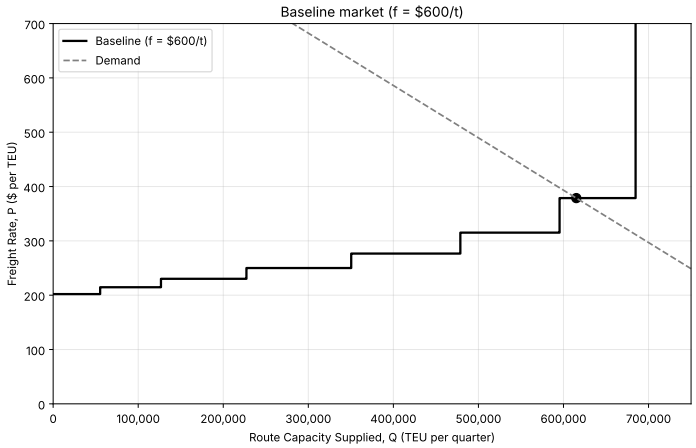


Figure 2: Marginal cost function with baseline fuel price

Plot of the marginal cost function with a set distribution of ships and at a baseline fuel price of $600 per ton, and demand is linear. The largest ships have the lowest marginal cost and go the fastest. Then smaller ships have higher marginal costs.

tity decrease, and the largest ships increase profits relative to the baseline. With the smallest ships leaving, the average ship size increases and frequency increases to compensate for slower ships. Finally, for large fuel price increases, there must be a proportional slowdown (f=$1500 in Figure 3). The smallest ships again earn negative profits. Now, however, the quantity decrease outweighs the price increase, and all ships earn lower profits than in the baseline. Again, with the smallest ships leaving, the average ship size should increase, but utilization concerns of the largest ships mean average ship size might decrease because larger ships cannot get filled. Frequency will still increase to compensate for slower ships.

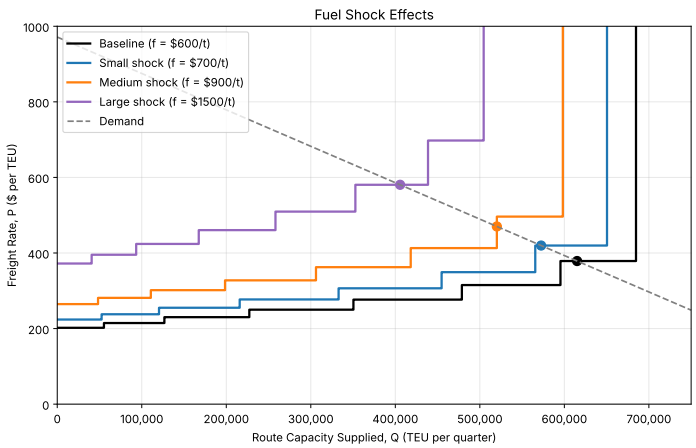


Figure 3: Marginal cost function with different fuel cost shocks

Plot of the marginal cost function with a set distribution of ships and at a baseline fuel price of $600 per ton, with linear demand and examples of higher fuel prices. f=$700/t represents a small increase in fuel price. f=$900/t represents a moderate increase in fuel price. f=$1500/t represents a large increase in fuel price.

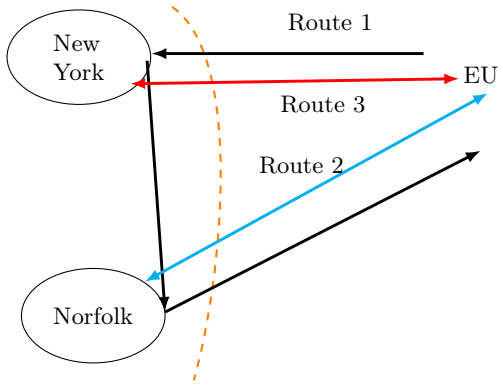


Figure 4: Example of Routes

Example routes include treated route (Route 1) and several untreated routes (Routes 2 & 3). The dotted line represents the ECA boundary that each route crosses. The arrow from New York to Norfolk for Route 1 is a fully treated segment.

3 Data

To estimate how the ECA changed container shipping supply characteristics, I construct a route-by-quarter panel from the US Army Corps of Engineers' Institute of Water Resources entrance and exit data (port calls) for all ships passing through US ports.

The entrance and exit data document every ship entering or leaving a US port, tracking: the day, ship identification, size of the ship, draft, the port of interaction, and port of origin or destination (depending on whether the ship is entering, origin, or exiting, destination). I filter for commercial dedicated container ships and specialty cargo ships with container carrying ability. I then construct the individual trip segments of a container ship moving along its route. With the port exits, I observe the start of any trip and its anticipated end. With port entrances, I observe the end point and its documented start. This provides data on the start and end of a trip segment, and I combine them to get an observed trip segment. As an example, for a trip segment from Newark, New Jersey, to Baltimore, Maryland, I observe the day a ship exits the port at Newark with Baltimore as the stated next destination. I also observe the day the same ship enters the port of Baltimore with its stated origin as Newark. I then match these two observations together as a single trip segment. From the trip segments, I calculate the travel time in days, regional origin, and destination for each segment. This provides the basis to track container ships taking trips along the US portion of container routes.

I define a complete trip of a container ship in the US by incorporating the start and end points of the US portions of container routes. I use the observed port entrance with an international port of origin (start) and observed port exit with an international port of destination (end). For example, I observe an entrance to New York, NY, with the stated origin as Rotterdam, Netherlands, and that would serve as the start of a trip. For the end, I observe an exit from Norfolk, VA with a documented destination being Freeport, Bahamas, and that would be the end of the US trip (Route 1, Figure 4). I then order the trip segments chronologically between the trip start and the trip end. I then mark the series

of trip segments, with the start and end, as one complete trip. For example, once I've sorted by date, I have the arrival in New York from Rotterdam, then my trip segment from New York to Norfolk, then my exit from Norfolk to Freeport. I connect this as one trip to the US with Rotterdam as the origin of the trip.

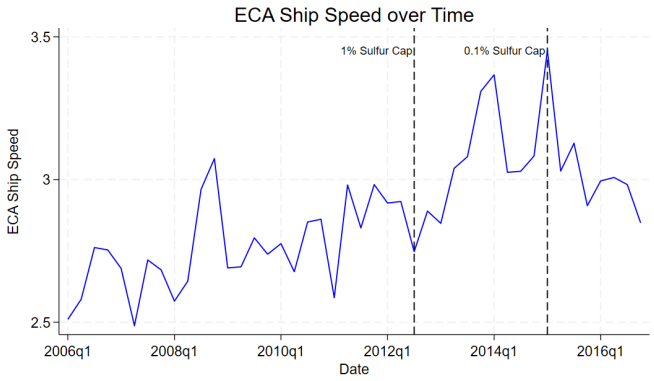
	Mean	SD	Min	Max	N
Outcomes					
Avg Ship Size (Tons/ship)	16364.05	21276.20	0.00	166410.00	19712
Frequency (count/quarter)	7.39	24.78	0.00	412.00	19712
Travel Time (days)	2.87	1.99	0.00	27.00	7326
Treatment					
Treated Routes	0.68	0.47	0.00	1.00	17380

Table 1: Summary Table

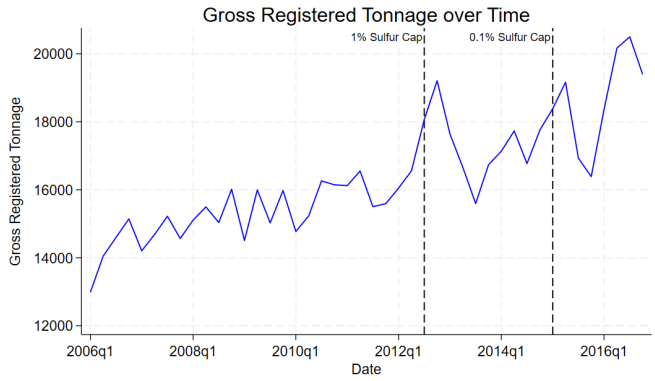
Note: The unit of observation is route-quarter. Summary statistics are reported for each outcome variable.

With the trips constructed, I identify the container route taken. I define a route as the continent a ship arrives from and each of the ports it stops at while in the US. Identifying the route for each trip to the US, I calculate the outcome variables and define treatment, aggregating to the route and quarter level. For average ship size, trip count, and average trip length, I average across all trips made in that quarter. For average speed, I average the time it takes to travel between each US port along a route; since the distance is fixed, the time it takes to travel that distance proxies for speed. For frequency, I sum the total number of trips for a quarter. I define a route as treated if it has multiple stops in the US. For example, the route from the EU to New York, NY, and then to Norfolk, VA is a treated route because it has a leg in the US (Route 1 in Figure 4). A route from the EU to New York and back is untreated (Route 2 in Figure 4). This results in 267 treated routes and 128 control routes.

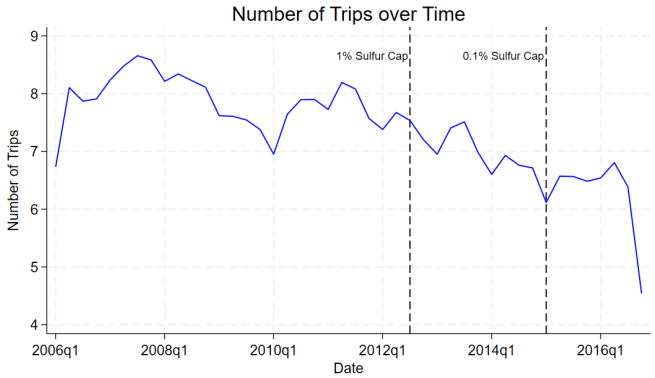
The aggregate statistics for the three outcomes of interest show that ships slowed down (Figure 5a) while in US waters and ships got larger (Figure 5b), following the model predictions. The frequency of trips goes down over time, indicating that if treated routes are seeing an increased frequency due to the ECA, it is offsetting the trend of fewer ships coming to the US rather than increasing the number of trips along a route (Figure 5c).



(a) Average Travel Time over Time



(b) Average Ship Size over Time



(c) Average Trips per Quarter over Time

Figure 5: Summary Statistics for each Outcome Variable

Note: The unit of observation is the average across the route-quarter panel.

4 Estimation and Results

4.1 Ship Size and Frequency

4.1.1 Estimation

To estimate how the ECA changed ship size and frequency, I use a difference-in-differences approach that compares economically similar routes, except that the treated routes have legs of the route completely within the ECA. Since all routes are exposed to the ECA, I use similar routes, but the treated routes have more than one stop in the US, while control routes have a single stop. The goal of the control group is to capture the counterfactual world where the ECA does not happen and, in doing so, control for the industry's general trends in ship size and demand. With the industry trends accounted for in the control group, to credibly estimate the impact, I assume the start of the ECA is exogenous, and that route treatment is exogenous. From those assumptions, I estimate the impact of the ECA using a difference-in-differences (Equation 6) approach for each of the outcomes of interest, y_{it}: ship size, number of port calls, number of different ships arriving at a port, and ship speed while traveling in US waters.

$$y_{it} = \beta \text{ ECA}_t \text{ x Exposure}_{it} * \omega_i + \text{Route}_i + \text{Date}_t + \text{Int Origin}_i \text{ x Route Group}_t + \text{US Coast}_i \text{ x Year}_t + \epsilon_{it} \quad (6)$$

Equation 6 is at the route, i, by quarter level, t. ECA_t is the post-2012 indicator for the start of the ECA. $Exposure_i$ is the dummy variable for ECA exposure for each route. To account for routes having heterogeneity in use, I weight by the number of total trips on that route, ω_i. Route fixed effects control for differences across routes consistent across the sample. Date fixed effects control for quarter-level time shocks. The Route Group by International Origin fixed effects allow for differences in trends for routes with large differences in use in the pre-period and from different international starting points. This pairs economically similar routes by matching ports with similar levels of shipping and from similar locations. I also include US coastal region by year fixed effects to allow for different time trends for the different coasts of the US. This helps to control for differences in changing regional US demand and transportation infrastructure.

4.1.2 Identification

The goal of the identification strategy is to compare treated and untreated routes with a similar population of ships and a similar demand for shipping to capture the non-ECA changes to the shipping market. The ideal untreated route would have the same sets of goods at the same level amounts, experiencing the same global shocks, and with the same

ships. To make that comparison, I include route and time fixed effects, Interaction of Region by Year, and International Origin and Route Size Categories, then I weight by the number of trips made over the whole dataset. The route fixed effects absorb the constant variations across routes, such as capacity constraints, either larger ones like the Panama Canal or smaller ones at the port level. Time fixed effects control for common shocks that affect all routes in a quarter. Examples include global changes in crude oil prices. A global crude oil price changes will affect demand for transportation and shipping decisions equally for all routes and absorbed by the time fixed effect.

The region by year interaction allows for different trends for different regions of the country. I created categories for the East Coast, West Coast, Gulf Coast, and Great Lakes. Different coastal regions have different economic activities and ship different goods to and from places. For example, the Gulf Coast specializes in petroleum products, while the East Coast receives consumer products. The coastal region by year allows those regions to have those different trends and shocks. Returning to crude oil, for example, the Gulf Coast might ship less because of a price drop since the region produces crude oil, while the East Coast could ship more because consumer products that use oil as an input are cheaper.

Route groups by international origin; interaction groups route by their usage and which continents the ships come from. This allows for different trends on routes with different usage levels from different continents. The international origin fixed effect absorbs the differences in economic conditions and goods traveling between different international origins. In container shipping the east-west routes of US-China, and US-Europe represent the main trading routes and the most heavily used and serve to bring goods to the US (Stopford, 2009). The north-south trading routes between the US and South America typically take goods south (Stopford, 2009). To allow for additional differences in usage across these routes, I include route group controls. A heavily used route from Europe is different from a less used route (Table 2). The interaction of the two fixed effects allows for different trends across these distinct categories.

This estimation approach could be undermined by ships traveling between treated and untreated routes because of the ECA, or by regulations coinciding with the ECA that have a similar effect. Ships switching between treated and control routes present a classic violation of difference-in-differences assumptions. To verify the lack of interaction, I look at the level of switching between treatment and control and find it stable across time (Figure 6). The limitation of this approach is that I cannot observe the goods themselves and cannot verify if there is an underlying shift in containers.

Regulations and behavior that could pose a threat to identification are those during the period of 2006 to 2016 that surrounded large fuel cost adjustments, such as slow steaming, and environmental regulations like the Energy Efficiency Design Index. Slow steaming started in response to high oil prices in 2007; it was implemented along major shipping routes, including the Europe to Singapore route, and was adopted globally by 2010 (Maersk, 2011). Since slow steaming started prior to 2012, if it affected treated and untreated shipping routes differently, the pre-trend estimates will show the effect being positive prior to the ECA. Energy Efficiency Design Index is a regulation from the International Maritime Organization (The United Nations organization that regulates the global maritime industry) that sets efficiency

Rank	Route	Total trips	Trips (pre-ECA)	Trips (post-ECA)	Avg GRT
1	Carib/Cent Amer -> MIAMI, FL	12842	8722	4120	5894
2	Carib/Cent Amer -> PORT EVERGLADES, FL	11518	7028	4490	9179
3	Asia -> LOS ANGELES, CA -> OAKLAND, CA	8663	4754	3909	69554
4	Asia -> LONG BEACH, CA -> OAKLAND, CA	5825	4089	1736	73482
5	Europe -> NEW YORK, NY -> NORFOLK, VA	4750	2960	1790	47169
6	Carib/Cent Amer -> HOUSTON, TX	4448	2427	2021	18474
7	Carib/Cent Amer -> HOUSTON, TX -> NEW ORLEANS, LA	3385	1970	1415	45856
8	Asia -> NEW YORK, NY -> NORFOLK, VA	2672	1382	1290	56776
9	Asia -> LONG BEACH, CA	2157	1221	936	67240
10	Asia -> LOS ANGELES, CA	2086	1632	454	70316
11	Carib/Cent Amer -> LOS ANGELES, CA -> OAKLAND, CA	1776	1208	568	50236
12	Carib/Cent Amer -> PANAMA CITY, FL	1685	1176	509	2719
13	Carib/Cent Amer -> FORT PIERCE, FL	1662	1649	13	996
14	South America -> HOUSTON, TX	1597	1076	521	7914
15	Asia -> SEATTLE, WA	1568	1182	386	63198
16	Carib/Cent Amer -> LOS ANGELES, CA	1526	938	588	31554
17	Carib/Cent Amer -> GULFPORT, MS	1525	1057	468	12863
18	Carib/Cent Amer -> SAVANNAH, GA -> NEW YORK, NY	1331	987	344	51803

Table 2: Ranking of Routes by Usage

Note: Ranking of Routes by usage for the top 18 routes. It shows usage pre- and post-ECA and average ship size along that route.

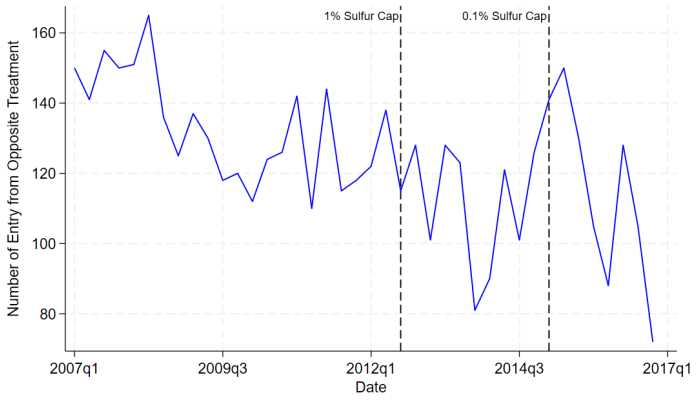


Figure 6: Number of Ships Switching between Treated and Untreated Routes per Quarter

Note: Unit of observation is route-by-quarter level. Counted for all routes in a quarter. 2006 excluded because there is no ability to tell what ships are switching.

standards on new orders for building ships starting January 2013 (Ghosh, 2022). If treated routes saw these more efficient ships, it could be evidence that the ECA had an effect by that response. The research design, however, does not capture that effect. Within the research design, fuel cost spikes impact more efficient ships less and would be less responsive to a fuel shock. Then more efficient ships on treated routes would bias the estimates downward. In addition, it takes 2 years to build a new container ship (QH Shipping, 2026). The building lag limits the impact of the EEDI to 2015 and 2016, but this is where the ECA impacts are being observed.

4.1.3 Results

I find that the ECA changed shipping supply characteristics. Along exposed routes, the ECA led companies to switch to larger ships and increase frequency, while each ship along the route operated at slower speeds. I start by showing the results of event studies for ship size and frequency. The event studies show evidence of the parallel trends assumption before the ECA, then show the effect evolving after the start.

Starting with ship size, I find that before the ECA, routes exposed to the ECA respond similarly to the control, indicating parallel trends holding (Figure 7). After the ECA, from 2012 to 2014, there is a small and statistically not significant response. In 2015 and 2016, after the tightening of the regulation, the point estimates are larger but still not statistically significant. This is likely due to limited statistical power for the annual estimates. The aggregate difference-in-differences results show that after the ECA, the average ship on a treated route was about 2400 gross tons larger (15% increase) than an average ship on a similar control route (Table 3).

A similar pattern is seen for the frequency of trips (number of trips) (Figure 8). Before the ECA, there is no statistical difference between the exposed routes and the control. After the ECA, the point estimates show an increase in the num-

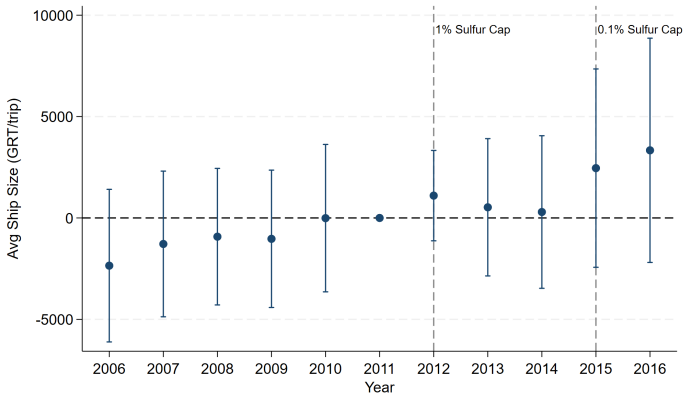


Figure 7: Event Study: Average Ship size relative to 2011 over time

Note: The unit of observation is route-quarter. The estimates represent a route's additional response because of the ECA. Robust standard errors are at the port and three-year increment level. All confidence intervals are reported at the 95% level.

	(1) Avg Ship Size (GRT/trip)	(2) Frequency (count/quarter)
ECA x Route Treat	2419.980**	26.940***
	(1215.246)	(10.127)
Route FE	Yes	Yes
Date FE	Yes	Yes
Interaction FEs	Yes	Yes
Observations	17424.000	17424.000
Mean Y Var	16358.541	7.370
Adjusted R^2	0.820	0.951

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Estimation Results

Note: The unit of observation is route-quarter. The estimates represent a route's additional response because of the ECA. Robust standard errors are at the route and three-year increment level. All confidence intervals are reported at the 95% level.

ber of trips for treated routes relative to untreated, however, no year individually is statistically significant. The aggregate difference-in-differences results for the number of trips show that after the ECA, the number of trips on a treated route was about 26 more than the number of trips on a similar control route (Table 3) and that result is statistically significant at the 1% level.

4.2 Travel Time

With evidence that ships get larger and frequency increases after the ECA, I now turn to the impact of the ECA on ship speed. Due to a limitation of the Army Corps data, there are no observations for the control group used in the difference-in-differences approach, so I follow the method of Borusyak et al. (2024) to estimate the impact of the ECA.

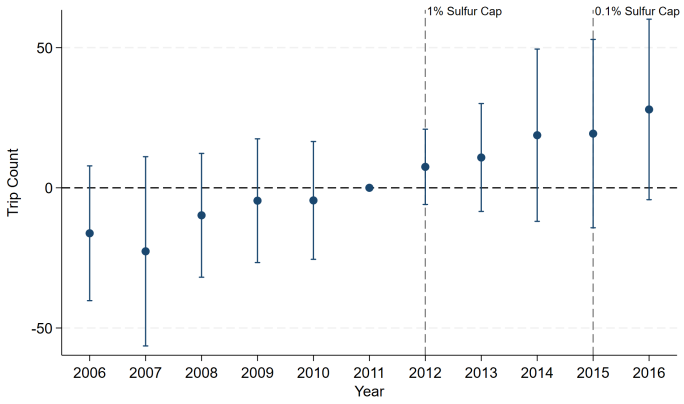


Figure 8: Event Study: Average number of trips on a route relative to 2011 over time

Note: The unit of observation is route-quarter. The estimates represent a route's additional response because of the ECA. Robust standard errors are at the route and three-year increment level. All confidence intervals are reported at the 95% level.

The results of the imputation procedure show that ships slowed down after the ECA, matching the model's predictions.

The Army Corps data is missing travel time information for the control group. I calculate travel time by taking the difference between the day a ship leaves a US port and the day it enters another US Port. This means control routes (routes with only one stop in the US) have no travel time information since there is no data on the start or end point, as one of those is not a US port. Given the missing control group, another approach is needed to estimate the impact in the absence of a clear counterfactual.

To estimate the impact of the ECA on speed, I use an imputation approach, following Borusyak et al. (2024), and using the stylized model to guide model selection. In the model, speed, measured as travel time over a fixed distance ($\textbf{Tvl. Time}_{it}$), varies with fuel costs ($\textbf{Crude}_t$), ship size ($\textbf{Size}_i$), and size rental prices ($\textbf{Capital}_t$), (Equation 7). The imputation approach first estimates the relationship between travel time and non-regulated fuel costs, cost of size, and ship size for the period before the ECA (2006 to 2011). From that equation, I estimate travel time ($\widehat{\textbf{Tvl Time}}$) for each trip segment for 2006 to 2016. This serves as the counterfactual travel time because it assumes the relationship between crude prices, ship size, and Treasury yields is the same pre- and post-ECA. To estimate the impact of the ECA, I take the difference between the observed and the counterfactual travel time. Then, I estimate how the difference between the observed and imputed travel time changes after the start of the ECA (Equation 8).

$$log(\textbf{Tvl Time}_{it}) = \alpha_i + \beta_1 log(\textbf{Crude}_t) + \beta_2 log(\textbf{Capital}_t) + \beta_3 log(\textbf{Size}_i) + \beta_4 \textbf{Route}_i \text{ x } \textbf{Stop}_i + \epsilon_{it} \tag{7}$$

To estimate the counterfactual, I use data on ship travel time for each trip, ship size for each trip, crude oil prices, and 10-year US Treasury yields from 2006 to 2016. The data on travel time and ship size come from the Army Corps data

	(1) Travel Time (Days)
Crude Prices (2 Month Lag/Log($))	0.005
	(0.008)
10-Yr Yields (log(%))	-0.125***
	(0.008)
Ship Size (log(GRT))	-0.041***
	(0.006)
Constant	1.242***
	(0.073)
Observations	33,038
Adjusted R^2	0.492

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Estimated Betas

Note: The unit of observation is ship-trip. The estimates represent the change in travel time with a change in each variable. Confidence intervals are reported at the 95% level.

set described previously. The raw data show that, on average, ships slowed down after the ECA (Figure 9). Here I use the trip-level observations to observe the time it takes to make any trip between two US ports. I use crude oil prices instead of high-sulfur shipping fuel because crude prices are available for the whole period. I use a lag of 2 months for prices to account for the time it takes for crude oil prices to pass to ships. The data on 10-year Treasury yields are a proxy for the rental price of size. Cost of capital rates are not readily available for individual shipping firms. However, container shipping companies borrow from banks against their whole fleet. Banks typically use a risk-free rate like US Treasury yields as the baseline to set borrowing costs for ships (Stopford, 2009). This results in a panel of trip segments, i, over time t, where I observe the trip an identified ship took, the size of the ship, the days the trip started and ended, how long the trip took, and the current price of crude oil and current Treasury yields. To compare travel times on the same segments of the route, I include a route and stop interaction (Equation 7). The average travel time for the panel is 2.8 days and fluctuates between 2.5 and 3 days between 2006 and 2011.

Estimating the baseline parameters for a ship's observed speed on a trip in the period before the ECA. I linearize the model by taking logs of each variable: travel time, Treasury yields, crude oil prices, and ship size (Equation 7). To capture that different trips are different lengths, I include a route and stop interaction fixed effect. The result of the estimation shows the betas match the model's predicted sign for ship size, Treasury yields, and the constant. As predicted, crude oil prices are positive, but the coefficient is not statistically significant.

Counterfactual travel times for treated routes are imputed using the coefficients estimated on untreated observations (Equation 7). Graphically, the difference between observed and imputed travel time is close to zero for the pre-period. After the ECA, the difference gets larger, showing the ships are going slower than what the imputation would predict

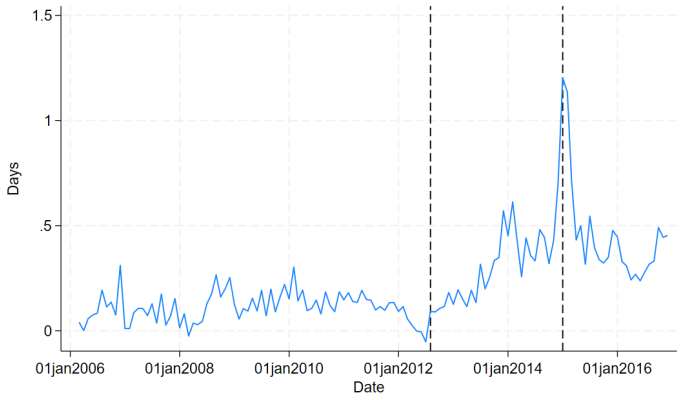


Figure 9: Average Difference between Observed Travel Time and Estimated Travel Time

Note: The data is averaged across the observations at the ship-trip level

(Figure 9). I estimate the effect by the change in difference pre- and post-ECA (Equation 8). Following the difference-in-differences approach, I weight by the number of trips seen along the route over the whole period. To capture the uncertainty in the estimates, I follow Jepsen et al. (2025) and bootstrap my standard errors with 10000 samples with replacement.

$$\textbf{Tvl. Time}_{it} - \textbf{Tvl. } \widehat{\textbf{Time}}_{it} = \beta \ \textbf{ECA}_t * \omega + \epsilon_{it} \tag{8}$$

The results show ships slowed down after the ECA. Before the ECA, there is no statistical difference between the observed and imputed travel times, giving evidence of the predictive power of the imputation approach (Figure 10). After the start of the ECA, the estimates show ships slow down. The average effect is a 16 percent slowdown in ships after the ECA (Table 5). The slowdown increased after 2012, peaking in 2015 and ending slightly lower in 2016.

5 Heterogeneity

The aggregate results show a response that follows the expected results for the moderate fuel cost shock: ships slow down, ship size increases, and ships travel the route more frequently. One conclusion of the model is that the magnitude of the cost change leads to different changes in ship size. For the aggregate results, the response was larger after the tightening of the ECA in 2015. Another way to examine the effect of the size of the fuel cost shock is by distance within the ECA. With the ECA, the cost change is heterogeneous because the distance traveled in the ECA changes route to route, and longer distances experience larger cost shocks.

Different cost shocks in the model have different responses. A small shock can be adapted to by all ships; each ship

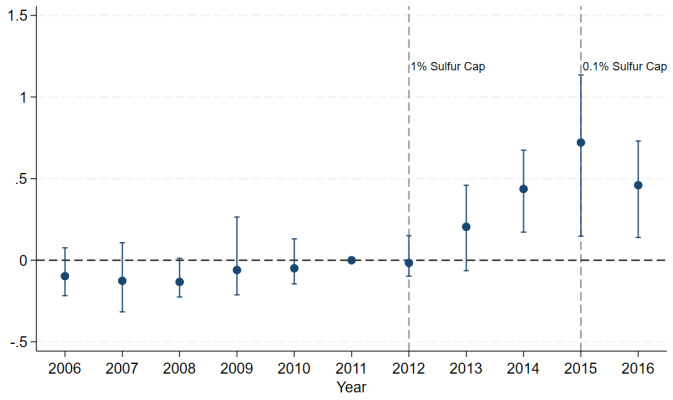


Figure 10: Results of Imputation Method

Note: The unit of observation is ship-trip. The estimates represent a route's additional response because of the ECA. Bootstrapped standard errors are at the route and three-year increment level. All confidence intervals are reported at the 95% level.

	(1) Speed (Imputed-Days)
Post 2012	0.479***
	(0.105)
Route FE	Yes
Date FE	No
Quarter of Year FE	No
Observations	52,552
Mean Y Var	0.319
Adjusted R^2	0.056
Wild bootstrap p-value	0.001
Wild bootstrap 95% CI	[0.200, 0.733]

Table 5: Aggregate Results of Ship Travel Time Imputation Approach

Note: The unit of observation is ship-trip. The estimates represent a route's additional response because of the ECA. Bootstrapped standard errors are at the route and three-year increment level. Confidence intervals are reported at the 95% level.

slows down to minimize fuel costs, and because the price might go up slightly as well, each ship can even earn more revenue than before. A moderate cost shock requires a larger drop in speed to optimize for higher fuel prices. The marginal ships earn below-zero profits because the quantity demanded gets reduced. The largest ships earn higher profits than the baseline because of their fuel efficiency and the new higher prices. A large fuel shock leads to dramatic reductions in quantity demanded and higher prices. Different from the moderate case, the demanded quantity reduction overwhelms the higher prices, and all profits are worse off than before. Given the large reduction in quantity demanded, it might be more difficult to fill larger ships, and they exit the market.

To see how distance might change the response, I measure the distance of the US sections of the ECA, then bin routes into separate bins for short, medium, and long routes, which proxy for small, moderate, and large fuel cost shocks. To

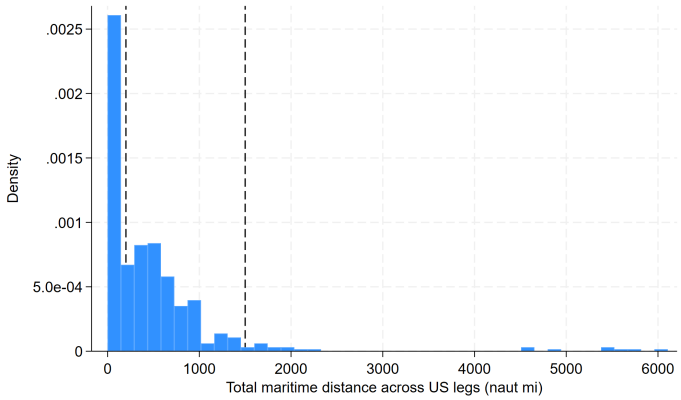


Figure 11: Distribution of route distances within the ECA

Note: Observations of route distances for the segments within the ECA. Dotted lines represent 200 nm and 1500 nm.

measure the distance between ports, I use the seaports package in Python [2]. This open-source data measures the distance between ports accounting for land masses that ships have to avoid. Plotting the length of the US sections of routes shows a large difference between the shortest routes at 41 nm and the longest at about 6000 nm (Figure 11). Most routes are below the average of about 1500nm, with very long outliers near several thousand nm.

I estimate the effect following the difference-in-differences approach as before, comparing the treated routes in each bin to the whole set of counterfactual routes.

The estimates for heterogeneity follow the expected results of the model (Table 6). Following the predictions of the model, frequency increases along each of the routes to maintain the level of service. Short routes experienced no statistical change in ship size compared to the counterfactual. Medium-length routes saw a statistically significant increase in ship size. The long routes saw a decrease in ship size.

6 Robustness

To show how the weighting choice and clustering choices affect the results, I re-ran the difference-in-differences regression without weighting by total trips and with clustering at different levels. For the clustering check, I first cluster at the route-by-year level and then at the route level. The unweighted regression have the same sign, but magnitude is lower for both ship size and frequency. The results are due to less-used routes being treated equally to more-used routes in estimation, and those less-used routes might be less able to adjust (Panel A: Table 7).

For clustering, we see that route-by-year clustering improves statistical precision because more observations are considered independent (Panel B: Table 7). For clustering at the route level, certainty goes down because the estimation

[2] https://pypi.org/project/searoute/

	(1) Avg Ship Size (DWT/trip)	(2) Frequnecy (count/quarter)
Panel A: Short Distance		
ECA x Route Treat	481.868	38.816***
	(2507.671)	(12.323)
Route FE	Yes	Yes
Date FE	Yes	Yes
Interaction FEs	Yes	Yes
Observations	6,996	6,996
Mean Y Var	11977.936	9.323
Adjusted R^2	0.872	0.959
Panel B: Med Distance		
ECA x Route Treat	2651.188**	26.910***
	(1260.466)	(10.268)
Route FE	Yes	Yes
Date FE	Yes	Yes
Interaction FEs	Yes	Yes
Observations	15,312	15,312
Mean Y Var	16370.285	8.559
Adjusted R^2	0.826	0.951
Panel C: Long Distance		
ECA x Route Treat	-5033.719*	26.110***
	(2695.928)	(9.023)
Route FE	Yes	Yes
Date FE	Yes	Yes
Interaction FEs	Yes	Yes
Observations	6,248	6,248
Mean Y Var	12151.300	10.000
Adjusted R^2	0.880	0.959

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Estimates of the Ship Size and Frequency for Routes of Different distances within the ECA

Note: The unit of observation is route-quarter. The estimates represent a route's additional response because of the ECA. Panel A shows the results for treated routes with a distance of less than 200 nm within the ECA. Panel B shows the results for treated routes with a distance of between 200 nm and 1500 nm within the ECA. Panel C shows the results for treated routes with a distance greater than 1500 nm within the ECA. Robust standard errors are at the route and three-year increment level. All confidence intervals are reported at the 95% level.

approach considers errors for all years correlated (Panel C: Table 7). This may be unrealistic because it considers the errors for 2016 observations to influence 2006. The results show frequency estimates are still significant at the ten percent level. Ship size is no longer significant even at the ten percent level.

7 Conclusion

Prior research focused on the short-run impact of the ECA, finding that there were pollution reductions and ships slowed down while in ECA waters. My results show the additional margins of adjustment made in response to the ECA. I also show that ships slowed down because of the ECA, but there was also an increase in the frequency of ships along exposed routes, and that the tightening of the ECA led to larger ships being used along exposed routes.

These additional margins of adjustment have implications for understanding and anticipating impacts from the ECA. One aspect if larger ships are serving the market, it means more container capacity. The estimated increased ship size of 2419 gross tons corresponds to about 318 containers per ship per quarter [3]. Using the elasticity of container capacity to ship price of -0.119 (Sui et al., 2026), you would anticipate a $5 decrease in the shipping price of a container. The EPA, in contrast, anticipated a $18 increase in the cost of shipping a container. This means those shipping goods along routes exposed to the ECA are better off because they are paying lower prices. For shipping companies, owners of larger ships will be better off because they earn higher profits despite the additional cost. The cost could be borne by the owners of smaller ships, who either earn lower profits or leave the market.

Finally, these results matter for port infrastructure investments because shipping becoming more concentrated and more frequent requires important adjustments for port logistics. Ports will need to handle more containers at one time and move to make space for the next ship. Research by Brancaccio et al. (2024) found optimal port investments should focus on large regional ports with large throughput due to returns to scale. The ECA further concentrating containers onto larger ships likely helps larger regional ports that can handle these ships more efficiently. This concentration means investment should be focused on those ports, but it also means benefits or profits flow to those larger ports.

[3] Formula from Abramowski et al. (2018): Ship Size (Gross Tons) $= -1097.11 + 11.049 * \text{TEU}$

	(1) Avg Ship Size (GRT/trip)	(2) Frequency (count/quarter)
Panel A: Port by 3 Year Level Cluster and unweighted		
ECA x Route Treat	1438.708*	1.989**
	(747.660)	(0.790)
Route FE	Yes	Yes
Date FE	Yes	Yes
Interaction FEs	Yes	Yes
Observations	17380.000	17380.000
Mean Y Var	16364.051	7.391
Adjusted R^2	0.443	0.923
Panel B: Route by Year Level Cluster		
ECA x Route Treat	2464.852**	27.394***
	(957.090)	(6.532)
Route FE	Yes	Yes
Date FE	Yes	Yes
Interaction FEs	Yes	Yes
Observations	17380.000	17380.000
Mean Y Var	16364.051	7.391
Adjusted R^2	0.820	0.951
Panel C: Route Level Cluster		
ECA x Route Treat	2464.852	27.394*
	(1760.017)	(15.331)
Route FE	Yes	Yes
Date FE	Yes	Yes
Interaction FEs	Yes	Yes
Observations	17380.000	17380.000
Mean Y Var	16364.051	7.391
Adjusted R^2	0.819	0.951

Table 7: Robustness Check Table

Note: The unit of observation is route by quarter. Panel A represents difference in difference results for unweighted regression clustered at the route and three year level. Panel B represents the weighted difference in difference regression with clustering at the route by year level. Panel C represents the weighted difference in difference regression with clustering at the route level. All confidence intervals are reported at the 95% level.

References

Abramowski, T., Cepowski, T., and Zvolenský, P. (2018). Determination of regression formulas for key design characteristics of container ships at preliminary design stage. *New Trends in Production Engineering*, 1:247–257.

Adland, R., Fonnes, G., Jia, H., Lampe, O. D., and Strandenes, S. P. (2017). The impact of regional environmental regulations on empirical vessel speeds. *Transportation Research Part D: Transport and Environment*, 53:37–49.

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

Beverelli, C., Benamara, H., Asariotis, R., and Secretariat, U. (2010). Oil prices and maritime freight rates :: an empirical investigation : technical report /: by the UNCTAD Secretariat.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, 91(6):3253–3285.

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2020). Geography, Transportation, and Endogenous Trade Costs. *Econometrica*, 88(2):657–691. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15455.

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2023). The impact of oil prices on world trade. *Review of International Economics*, 31(2):444–463. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/roie.12632.

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2024). Investment in infrastructure and trade: The case of ports. Working Paper 32503, National Bureau of Economic Research.

Cullinane, K. and Khanna, M. (1999). Economies of scale in large container ships. *Journal of Transport Economics and Policy*, 33(2):185–207.

Fowlie, M. (2010). Emissions trading, electricity restructuring, and investment in pollution abatement. *American Economic Review*, 100(3):837–69.

Fowlie, M., Reguant, M., and Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1):249–302.

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

Ghosh, S. (2022). What is energy efficiency design index? Marine Insight. Last updated 1 May 2024. Accessed 22 October 2026.

Hansen-Lewis, J. and Marcus, M. (2025). Uncharted waters: Effects of maritime emission regulation. *American Economic Journal: Economic Policy*, 17(1):37–69.

Hummels, D. (2007). Transportation Costs and International Trade in the Second Era of Globalization. *Journal of Economic Perspectives*, 21(3):131–154.

International Maritime Organization (2014). Greenhouse gas studies 2014. Accessed: 2024-09-21.

International Trade Administration (2024). Maritime services trade data. https://www.trade.gov/maritime-services-trade-data. U.S. Department of Commerce. Last modified October 1, 2024. Accessed May 4, 2026.

Jepsen, C., Mueser, P., Troske, K., and Jeon, K.-S. (2025). The Benefits of Alternatives to Conventional College: Comparing the Labor-Market Returns to For-Profit Schools and Community Colleges. *Journal of Human Resources*.

Kalouptsidi, M. (2014). Time to build and fluctuations in bulk shipping. *American Economic Review*, 104(2):564–608.

Klotz, R. and Berazneva, J. (2022). Local standards, behavioral adjustments, and welfare: Evaluating california's ocean-going vessel fuel rule. *Journal of the Association of Environmental and Resource Economists*, 9(3):383–424.

Klotz, R. and Berazneva, J. (2025). Correlated pollutants, avoidance, and local environmental policy. *The Energy Journal*, 0(0):01956574251330842.

Lan, X., Tao, Q., and Wu, X. (2023). Liner-shipping network design with emission control areas: A real case study. *Sustainability (Switzerland)*, 15(4).

Maersk (2011). Slow steaming – the full story. http://www.maersk.com/Innovation/WorkingWithInnovation/Documents/Slow%20Steaming%20-%20the%20full%20story.pdf. Archived 2 November 2012 at the Wayback Machine: https://web.archive.org/web/20121102073152/http://www.maersk.com/Innovation/WorkingWithInnovation/Documents/Slow%20Steaming%20-%20the%20full%20story.pdf. Accessed 22 October 2026.

Mirza, D. and Zitouna, H. (2009). Oil Prices, Geography and Endogenous Regionalism: Too Much Ado About (Almost) Nothing. *CEPII research center, Working Papers*.

Moore, T. J., Redfern, J. V., Carver, M., Hastings, S., Adams, J. D., and Silber, G. K. (2018). Exploring ship traffic variability off california. *Ocean & Coastal Management*, 163:515–527.

Notteboom, T. E. and Vernimmen, B. (2009). The effect of high fuel costs on liner service configuration in container shipping. *Journal of Transport Geography*, 17(5):325–337.

Peters, A. (2024). Beached assets? Capital turnover and emissions in shipping. Job Market Paper.

QH Shipping (2026). How long does it take to build a ship? https://www.qh-shipping.com/How-Long-Does-It-Take-to-Build-a-Ship-id00046135.html. Accessed 22 October 2026.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061.

Stopford, M. (2009). *Maritime Economics*. Routledge, London and New York, 3rd edition.

Sui, C., Wang, S., Liang, J., and Nomikos, N. K. (2026). Operating capacity, pricing and supply elasticity in container shipping markets. *International Journal of Finance & Economics*, 31(3):3139–3162.

Toyama, Y. (2024). Dynamic incentives and permit market equilibrium in cap-and-trade regulation. *American Economic Journal: Microeconomics*, 16(3):374–423.

U.S. Coast Guard Office of Commercial Vessel Compliance (2014). Frequently asked questions: North american emission control area (ECA). Technical report, United States Coast Guard. Revision: 12/3/2014. Originally published 8/30/2012.

US EPA (2009). Proposal to Designate an Emission Control Area for Nitrogen Oxides, Sulfur Oxides and Particulate Matter. Regulatory Proposal MPEC 59/6/5, International Maritime Organization.

US EPA (2010). Proposal to Designate an Emission Control Area for Nitrogen Oxides, Sulfur Oxides and Particulate Matter. *Assessment and Standards Division Office of Transportation and Air Quality U.S. Environmental Protection Agency*.

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

Zavitsas, K., Zis, T., and Bell, M. G. H. (2018). The impact of flexible environmental policy on maritime supply chain resilience. *Transport Policy*, 72:116–128.

Zhu, J. and Wang, J. (2021). The effects of fuel content regulation at ports on regional pollution and shipping industry. *Journal of Environmental Economics and Management*, 106.