Essays on the Industrial Organization of Ocean Shipping

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Samuel M. Bailey

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy

Thomas J. Holmes, Advisor

May, 2021

© Samuel M. Bailey 2021
ALL RIGHTS RESERVED

Acknowledgements

I am grateful to my advisor, Thomas J. Holmes, and committee members Doireann Fitzgerald, Amil Petrin, and Joel Waldfogel. All have provided excellent advice. Tom, Amil, and Joel have also been invaluable as members of the Minnesota Applied Micro Workshop. I would also like to thank the rest of the participants of the workshop for their comments and support. Serafin Grundl, Cailin Slattery and Helen Willis Banga all gave helpful feedback on early drafts of the first chapter. Finally, I would like to thank employees at the port authorities in Long Beach, Oakland, Houston, and Seattle for providing me with their data.

Dedication

To my parents, Michael and Suzanne Bailey.

Abstract

This dissertation consists of two chapters on the industrial organization of ocean shipping. In the first chapter, I study the investment decisions of US ports. Transportation infrastructure is characterized by two opposing forces: economies of scale that encourage centralization, and spread-out consumers that encourage dispersion. These forces may not be correctly balanced in the United States as decisions are made by many different regional authorities which receive large federal subsidies. I study seaports during a period when those on the East Coast were making investments to prepare for the larger vessels that could navigate an expanded Panama Canal. With data on all container imports and capital costs of major US ports, I estimate a model of the investment game that port authorities play. Competing ports invest more than a social planner would, even allowing for deviations from profit maximization, because they do not internalize their business stealing effects on others. In particular, the $1.7 billion expansion of the Port of New York and New Jersey would not have been chosen by a coastal authority. Social surplus would be over a billion dollars higher with coordination, the equivalent of about one year's worth of revenue for all the East Coast ports. Lowering federal subsidies appears to lower much of the gap. In the second chapter, I show how port productivity changed after a historic labor agreement. Across many industries, employers and workers often argue over technology adoption. In 2008, the International Longshore and Warehouse Union signed a contract agreeing that all ports on the US West Coast could fully automate their terminals, recognizing there would be job losses. Using a new, ship-level dataset of the labor it takes to unload ships, I study changes in productivity after the contract was signed and after one port automated. Having the ship-level data is important, as I show more aggregate measures would produce misleading estimates. I find productivity increased about 25% as a result of the new contract and an additional 15% among the port that actually automated. I find that the effects did not completely persist through the 2015 contract, even though the automation clause did not change, and suggest possible ways employee-employer relations may alter outside the written contract.

Contents

List of Tables

List of Figures

Chapter 1

Competition and Coordination in Infrastructure: Port Authorities' Decisions to Become "Big Ship Ready"

1.1 Introduction

In the United States, regional authorities often compete with one another. There has been a long line of research since at least Charles Tiebout in 1956 studying when this competition is more efficient than a centralized government. When businesses and residents have heterogenous costs and preferences, authorities that differentiate themselves can be beneficial. However, there may also be Prisoners' Dilemma scenarios, where both regions would benefit by committing to a lower level of investment. These investments are frequently subsidized by outside authorities, so lower levels may be beneficial not only to locals but the entire country.

Over the past ten years there has been a flurry of infrastructure investment by port authorities on the East Coast in response to the expansion of the Panama Canal, announced in 2006 and completed in 2016. The Canal grew to allow a ship size almost three times the pre-expansion maximum. Larger ships can carry more containers with the same size crew and proportionally less fuel, so there are increasing returns to scale. However, these economies of scale are wasted if a port is not deep enough to allow the ship to unload. These authorities are all regional and they rushed to invest to make their ports more attractive to importers potentially without internalizing the effects their investments would have on the business of other ports. These projects are also highly subsidized by the federal government, making the investments even more attractive to individual ports.

In this chapter, I study whether competition between port authorities and the subsidies they receive have led to excess or insufficient investment in container ports. I find that there was at least a billion dollars in excess investment. By "excess investment" I mean investment that does not increase social welfare more than it costs. I estimate substitution patterns between ports using the universe of waterborne container imports and a hand-collected panel of port characteristics, which includes the money spent on investments to prepare ports for larger ships possible after the Panama Canal expanded. Using these substitution parameters and economies of scale parameters estimated from other work, I model the investment game port authorities play. These data are detailed enough that I can generalize the port authorities' objectives beyond that of profit-maximizers, and so my welfare results take seriously what the port authorities are actually striving for.

Though there has been a long literature studying regional competition, starting with (Tiebout (1956)). More recently, researchers have started using empirical techniques developed in industrial organization or auction theory to study fiscal competition, such as (Slattery (2019)). However, there has been relatively little work focusing on infrastructure investment. This is surprising, given 1) infrastructure spending is a large part of state total spending and 2) most infrastructure spending in the US is done on the regional level. For example, firm-level tax incentives, which have been the subject of much research, totaled a little under $50 billion in 2015 (Bartik (2019)). By comparison, state and local governments spent around $340 billion on transportation and water infrastructure in 2017. This is a large amount in dollar amount but also as a fraction of total infrastructure spending. In the same year, federal spending was below $100 billion, less than one-third that of states and localities (Office (2018)).

I focus on container ports. These are a vital part of US international trade: about half of trade by value is waterborne, and about half of waterborne trade is containerized (Chambers and Liu (2012)). There is a large amount of investment in these ports. In a 2012 survey, ports reported plans to spend $18.3 billion dollars over the next four years, and $22.6 billion from 2016-2020 (American Association of Port Authorities (2016)). All of these spending decisions are made by authorities often responsible to governments as small as municipalities. They may even be neighboring: Los Angeles and Long Beach are controlled by different port authorities, despite being less than five miles apart and in the same bay.

There are theoretical reasons competition could be either good or bad for total welfare. The reasons why competition could be good are intuitive: without it, ports may not be compelled to make investments that lower costs.[1] There may also be gains from variety, especially as regards to the different physical locations of each port. Importers consider not only port fees but the total freight charges on land and on sea. Different ports will

[1] For example, (Holmes and Schmitz (2001)) describe how port workers in the nineteenth century took advantage of their "bottleneck" position until development of rail forced them to improve.

be associated with different charges for different importers. These are positive outcomes of ports being decentralized. However, a long-standing result in the industrial organization literature is that in markets where a single firm's output decision affects the decision of other firms (i.e., in markets that are not perfectly competitive), there may be excess entry (Mankiw and Whinston (1986)). The reasoning is as follows: when a firm chooses to enter, it cares only about its expected payoffs. Upon entry, though, all other firms will reduce their output (or imports in the case of ports). This reduction in output is a cost not internalized by the firm, and as with any negative externality, there will be too much of it in a market equilibrium. The gains from variety are also not internalized by the firm but are a positive externality. Whether there is over, under, or optimal entry depends on which of these effects dominate.

In order to say whether there is over- or underinvestment, I need an efficient benchmark for what "optimal investment" looks like, and for this I need a model of the main agents involved: importers, carriers, and especially port authorities. Importers are located in particular places in the U.S. and ship goods from particular places abroad. They choose which port to have their good sent to, taking as given characteristics of the port (including the port's distance to the importer). Given their quasi-public nature, authorities may care about more than just profits, and so I allow for them to be concerned with total quantities in addition to profits. They compete with one another by investing in harbors that can accommodate larger ships, and these investments are subsidized by the federal government. Because larger ships have lower average costs, ports increasing their capacity is the equivalent of those ports shortening the distances between them and foreign origins. Thus though these investments are costly, they lower the effective price of importers who choose them.

I can identify these parameters from ship charter rates and technical specifications, the universe of waterborne container imports, and hand-collected data on port infrastructure and cost. A wide range of ship sizes existed even before the Canal expansion for routes that do not involve the United States, and a liquid and competitive charter market allows me to identify scale parameters by comparing the relationship between charter prices and size. For importers' substitution parameters, I use the fact that import origin and destination are exogenous. I can identify the cost of land transport from importers that have the same origin and same sized ship but different destinations. Ocean costs are identified similarly, from importers with same sized ship and same final destination but different origins. Finally, I take a revealed preference approach to bound the weight ports put on quantity. The scale and demand parameters allow me to predict what profits would be if ports did not invest. I therefore observe the realized profits, alternative profits, and cost of investment. The alternative profits generally exceed the realized ones, and this difference identifies the lower bound for the weight on quantity.

I estimate economies of scale and substitution parameters and then use these to consider counterfactuals. I find economies of scale are quite high, with average cost elasticities of 0.2-0.33.[2] Even for smaller vessels, the relative costs of shipping by land compared to ocean are much higher. I also find land costs are relatively high compared to sea: moving one container on land one kilometer is about 13 times as much as moving that same container one kilometer at sea on a 5,000 TEU vessel. Importers thus benefit from differentiation (in the form of geographically spread out ports) as well as economies of scale.

To see whether differentiation or economies of scale dominates, I consider counterfactual scenarios where not all ports invest in expanding. I find that ports coordinating among themselves would choose not to expand New York or Houston (depending on the exact form of the port authority objective). These choices provide greater total surplus for ports; they also provide greater total surplus as the cost savings from investment more than makes up for lower importer surplus. Although the port decisions without federal subsidies are still inefficient, the taxes that would be necessary to change behavior are much smaller. The subsidies have a large effect on the total welfare if not the physical allocation.

Importantly, New York would individually *not* want to go along with this plan; it would still be individually more profitable to invest. Just as tax competition between states can lower total surplus, it seems that ports competing with one another leads them to invest more than is socially desirable.

1.2 Literature

Many researchers have studied the effects of transportation infrastructure on total output or productivity; see (Turner et al. (2020)) and (Ramey (2020)) for recent overviews. There is no clear consensus on how important it is for output, but most acknowledge that the efficiency of transportation infrastructure varies greatly over space and time. (Turner et al. (2020)) find social welfare weights for several different modes of transportation and compare these to the implicit weights from government spending. They find the two sets of weights are wildly different, meaning infrastructure investment is inefficient. However, this and other papers study optimal national policy when states are not strategic agents. Outcomes that may seem inexplicable when there is a single planner, such as excessive investment in some areas, make much more sense when we consider that the investing agents are sometimes working at cross-purposes with one another. The paper closest to mine in this strand is (Hulten and Schwab (1997)), who study the effects of tax exemption on municipal

[2] That is, increasing ship size by 1 percent lowers average cost per container by 0.2-33 percent. A 8,000 twenty-foot equivalent unit (TEU) ship has average costs 12 to 20 percent lower than a 5,000 TEU one.

bonds and infrastructure subsidies. Because they are looking at aggregate investments, though, they cannot model the oligopolistic game being played by the agents as I do with port authorities.

There is another literature that does study competition between states or localities, starting with(Tiebout (1956)). The way states compete in these papers is through fiscal policy (such as tax credits or subsidies) or through public goods consumed directly by households (such as better schools or parks). A recent example is (Slattery (2019)), which studies states competing for firms locating there through subsidies. Such competition may be welfare enhancing if there are spillovers, as firms do not internalize the full surplus they generate. My work differs in that I study capital investment rather than fiscal policies. These policies are a much larger fraction of state budgets, about three times as much, as business incentives.

I also draw on results on "excess entry" in the industrial organization literature. A long-standing result is that in markets where a single firm's output decision affects the decision of other firms (i.e., in markets that are not perfectly competitive), there may be excess entry (Mankiw and Whinston (1986)). The reasoning is as follows: when a firm chooses to enter, it cares only about its expected payoffs. Upon entry, though, all other firms will reduce their output (or imports in the case of ports). This reduction in output is a cost not internalized by the firm, and as with any negative externality, there will be too much of it in a market equilibrium. The gains from variety are also not internalized by the firm but are a positive externality. Whether there is over, under, or optimal entry depends on which of these effects dominate. This has been studied theoretically by (Mankiw and Whinston (1986)) and found to be empirically relevant by (Berry and Waldfogel (1999)). Instead of profit maximizers, I study quasi-public authorities who may internalize some of these externalities already. I thus need to generalize the agents' objectives to allow for this fact.

Finally, many researchers in maritime economics have studied competition between ports. (De Langen and Pallis (2006)) and (Lee and Meng (2015)) discuss the growing devolution of port authority in Europe and Asia, respectively, and how the distribution of trade has changed over time. There has been less work studying competition among U.S. ports, possibly because there has always been less of a country-level port policy in that country. An exception is a Federal Maritime Commission report from 2012, though that focuses on competition between U.S. ports and Canadian and Mexican ports, not competition between U.S. ports themselves. (Ishii et al. (2013)) is a recent example that is closest to my approach. The authors model ports as making capacity investments in alternating periods, and then given those capacities, setting prices simultaneously. They derive several propositions whose results they compare to the ports of Busan and Kobe, but they do not do any estimation. My model of port investment is very similar to theirs; however, the bill of ladings microdata allows me to actually estimate the model, rather than derive only qualitative results.

1.3 Background

1.3.1 Port organization

Unlike many other countries, there is no single agency that administers U.S. seaports. Instead, they are governed by a variety of state, municipal, and regional authorities. Though all these authorities are ultimately responsible to elected officials, their operational independence varies: for example, the Massachusetts Port Authority is explicitly chartered to *not* be subject to control by other agencies, whereas the North Carolina State Ports Authority is within the state's transportation department (Sherman (2008)). In some cases, the authority board of commissioners is required to have at least some members from parts of the state far from the port, to guarantee the welfare of more than just the port city is being considered (Sherman (2008)). There are also differences in funding: most of the larger ports can cover their own operational costs, but may receive funding for select capital projects from the state or local governments. Others may have a guaranteed revenue from the state coming out of a specific tax or trust. The Virginia Port Authority, for example, receives revenue from the Commonwealth Port Fund, which is tied to highway taxes (Joint Legislative Audit and Review Commission (2013)).

Outside of government revenue or bonds, port authority's revenue comes from two major sources: rentals and the fees charged to ships that enter on its own terminals. There are many different kinds of fees for things like crane usage, refueling, water usage, etc. The main ones are wharfage, which is a per-weight or per-container fee for unloading, and dockage, which is a per-ship fee for docking at the port (though the fee will usually vary depending on the size of the ship) (American Association of Port Authorities (2019)). From the information available, wharfage is the larger of the two. For example, out of the approximately $398 million the Port of Los Angeles made from shipping services in 2017, $369 million, over 90%, were from wharfage. Thinking of fees as per container rather than per ship or per importer is therefore reasonable. I go into more details in the Data section.

Though ports are in the business of other types of imports, here I restrict attention to containers for three reasons. First, for most of the largest ports, containers are the primary source of revenue from shipping services. For the Port of Los Angeles, about 93% of imports by weight were from containers, and even more by revenue (Finance and of Port of Los Angeles (2019)). Second, containers are, by design, standardized. It is rare for two containers of the same volume to be charged different rates, regardless of commodity or even weight. (In theory, heavier containers may be charged more, but it is rare for weight rather than volume to be the limiting factor (Holmes and Singer (2018)).) Finally, because containers are often not opened until near the end of their journey, the port of entry (where it goes through customs) may be different from the port of unlading (where it comes off the ship). This allows me to see from customs data alone not only where a container physically entered the United States, but where it was unpacked.

1.3.2 Panama Canal expansion and container shipping

In 2006, Panamanians voted to add two new locks and to deepen existing channels. The decision was due in part to the growing size of containerships. As late as the mid-1990s, there was no containership in the world that exceeded the roughly 5,000 twenty-foot equivalent units[3] (TEU) of the Panama Canal. Around a decade later, when the plans for expansion began, the average size of a newly built ship was still below the limit. That was not the case when the expansion was finished, on June 26, 2016. By that point, the largest ships in the world were almost 20,000 TEU, and the average newly built ship was around 7,500 – well above the previous maximum allowable size. (Merk et al. (2015)). The new maximum of about 13,000 TEU was more than twice as large as the previous maximum.

There would have been no reason for the Canal to expand if it were not for the large economies of scale that exist for container ships. Average cost is almost always decreasing, limited only by demand and the infrastructure of canals and ports. There is "engineering" evidence of this: the resistance of water increases with the surface area of the vessel, not the volume. We can also see the economies of scale from the behavior of the carriers. Almost immediately after the expansion, carriers that had weekly services of two vessels around the maximum capacity switch to weekly service of only one larger vessel. Overall, the average vessel size of shipments from East Asia to the U.S. East Coast (which often though do not always go through the Panama Canal) increased by around 2,000 TEU following the expansion.

In the context of port infrastructure, the Canal expansion is important because it spurred ports on the East Coast to invest in becoming "Big Ship Ready." These investments were dramatic. The Bayonne Bridge, one of the largest steel arch bridges in the world, was raised over 50 feet at a cost of $1.7 billion in order to allow larger ships to enter the Port of New York and New Jersey. The Port of Savannah is currently near the end of a dredging project that deepened the harbor by 5 feet and cost over $970 million (Ronan (2020)). These projects all began after the decision to expand the Panama Canal was announced, and were explicitly presented as necessary because of the expansion. The ports also recognized that their investments would pay off potentially at the expense of other ports (Booth (2013)).

[3] One twenty-foot equivalent unit (TEU) is one half a standard 40-foot container.

1.4 Model

To say whether investment is excessive or insufficient, I need a way to compare the costs to the benefits it brings importers and ports. There is a trade-off between the economies of scale from concentrating shipping and the benefits of variety. To accurately measure the benefits I build on standard discrete choice demand models, incorporating important, individual-level features of this market: the economies of scale in ocean shipping, ocean distance, land distance, and port fee at the container level. For the costs I model the ports as choosing a level of investment that determines the largest ships that can unload there. Investment is expensive but indirectly lowers the cost for importers, bringing them to the port. Port authorities care about ships entering ports because it increases profits, but they also care about the total quantity for its own sake.

1.4.1 Overview

Importers have inelastic demand for containers from a foreign origin port o in U.S. destination d. Origin ports are always container seaports; the U.S. locations may be a port or an inland location. Importers choose a seaport j. This determines the journey a shipment takes from origin o to port j by water, and then from port j to destination d by land. There is a trade-off between ports that are expensive to get to by land but which may have low prices for the ocean leg of the journey. The total number of containers going from origin o to seaport (not final destination) j is $q_{o,j}$. This last number will be important in determining the size of the ship used.

Ocean carriers move containers from foreign ports to domestic ones. They take as given the total number of containers being shipped on the route $q_{o,j}$, and choose the size of the ship to minimize average cost.

Land carriers are a competitive sector that charge a constant price per unit distance.

Ports behave strategically, taking as given the actions of other ports. Their main decision is whether to invest in infrastructure that allows post-Panamax vessels to unload there.

1.4.2 Importers

I assume importers make decisions separately for each shipment[4] . They consider the price of moving the container by land from the domestic port to the final destination, $p_{land,j,d}$ and the price of moving it by sea from the foreign port to the domestic, $p_{sea,o,j}$. I allow each port to have an unobserved cost for all importers, ξ_j, that can vary over time. This may be how quickly the port can unload ships, or an especially valuable rail connection. Finally,

[4] A "shipment" may be one or more or even a fraction of a container. I treat all containers in the same shipment as having the same idiosyncratic cost ε_{ij}.

importers bear some individual cost shock, ε_{ij}, that is independent and uncorrelated with any other port feature.
The importers face the problem

$$\max_{j \in J} -\alpha \left(p_{sea,o,j} + p_{land,j,d} \right) + \xi_j + \varepsilon_{ij} \tag{1.1}$$

1.4.3 Ocean carriers

The ocean carriers' "problem" in the model is almost trivial. In reality, ocean carriers are strategic and important agents, and so this section of the model is the most simplified. Below I go into detail on which of these details may be important and why I am nevertheless comfortable making them for the model.
Ocean carriers take the quantities going from foreign ports to domestic ones as given. They have an increasing returns to scale technology and choose the ship size S that minimizes average cost, subject to the constraints of canals or ports. I model the choice of vessel size in a reduced-form manner as a function of the total route quantity, and domestic port capacity. (I assume foreign ports are not a binding constraint.) Let the vessel size be S and $S_{o,j} = S\left(q_{o,j}, \omega_j\right)$, where ω_j is the "maximum" allowed at port j. (Because investment is a binary decision, $\omega_j \in \{0, 1\}$.) $S\left(\cdot, \cdot, \cdot\right)$ is increasing in all three terms. No ship is "too large" to enter any port; however, it is very costly for larger ships to enter ports that are not Big Ship Ready. (This last feature is motivated by even "too large" ships being able to enter most ports at great inconvenience, see the Data section for examples.)
In treating vessel size as an exogenous function of only $q_{o,j}$ and ω_j, the model is missing the effects of traffic to nearby ports. For example, ships with only a few containers going to Long Beach may still be quite large because they make large deliveries to Los Angeles next door. However, the shape and direction of the function still holds: larger $q_{o,j}$ and larger ω_j both weakly increase the size.
After choosing the vessel that minimizes average cost, carriers simply charge this cost to the importers. In reality, the ocean carrier market is itself imperfectly competitive, and they likely charge a markup over their cost. For simplicity I abstract away from any strategic behavior on the carriers' side. As long as the price moves in the same direction as the costs, the direction of my results should still hold, and the magnitudes should not be too far off if markups are reasonably constant.
The carriers charge

$$p_{sea,o,j} = f\left(S\left(q_{o,j}, \omega_j\right), D\left(o, j\right); \beta_{sea}\right) + p_j + \varphi\left(S\left(q_{o,j}, \omega_j\right), \omega_j\right) \tag{1.2}$$

where p_j is the per-container port fee charged by port j, $D\left(o, j\right)$ is the distance from origin o to port j, and $\varphi\left(\cdot, \cdot\right)$ is a function of ship size and port size that is weakly decreasing in its second term (a larger port does not raise costs for anyone) but may be increasing in the first term if the port is not large enough. f is the average cost of moving a container on a ship of size $S_{o,j}$ the distance $D(o,j)$ with parameter(s) β_{sea}. In estimation, I parameterize this function based on work studying the economies of scale in ocean shipping.

1.4.4 Land carriers

Land carriers are a competitive industry with a constant marginal cost. They charge marginal cost per unit distance β_{land}. Thus $p_{land,j,d} = \beta_{land} D(j,d)$.

1.4.5 Port authorities

Port authorities play a game with each other in two stages: first, they choose whether or not to be Big Ship Ready, ω_j. Then, they set prices to maximize a weighted average of profits and total quantity, W_j. I explain the steps in reverse order.

Stage Two

The per-period payoff is

$$W_j = \left[\left(p_j - mc_j\right) q_j\left(\Omega, a_j, a_{-j}, p_j, p_{-j}\right)\right] + \theta_j q_j\left(\Omega, a_j, a_{-j}, p_j, p_{-j}\right) \tag{1.3}$$

Port authorities choose the price p_j that maximizes this value, taking as given the prices of other ports, p_{-j} the investment strategies, a_j and a_{-j}, and the state of other ports' preparedness, Ω. The first term on the right-hand side is standard, these are just the profits. The second term multiplied by θ_j is quantity alone, and θ_j is the weight port authority j places on it. Ports are quasi-public entities, and in their public statements often claim to be helping the entire regional economy. If they are purely profit-maximizing this term θ_j will be zero, but if they consider not only their own profits but their employment of workers, or business provided for local firms, it will be positive.[5]

Stage One

Taking the second stage objective functions as given, in the first stage ports decide whether or not to invest. They trade off the cost of investing with the higher q_j of more importers going through.

[5] I am agnostic about where θ comes from. One possible microfoundation is the following: the quantity unloaded, is a linear function of labor, $q_j = A_j \ell_j$. Then profits will be $p_j q_j - \frac{w}{A}_j q_j$, where w_j is the wage. Suppose that after choosing price, the port knows the surplus created will be bargained over, and therefore the wage paid will be a weighted average of the workers' outside options and the surplus $w_j = (1-\beta)Ap_j + \beta AU_j$, where β is the bargaining weight and U is the outside option. Thus when choosing price the port will internalize some of the wages, even though they appear in the financial reports as costs.

The ports authorities solve

$$\max_{a\in\{0,1\}}\left\{\tilde{W}\left(\Omega,a,a_{-j},\tilde{p},p_{-j}\right)-\mathbb{1}_{a=1}\left(1-\sigma_j\right)F_j\right\} \tag{1.4}$$

where

$$\tilde{W}\left(\Omega,a,a_{-j},p,p_{-j}\right)=\max_{p}W\left(\Omega,a,a_{-j},p,p_{-j}\right)$$

and $\tilde{p}=\arg\max_p W\left(\Omega,a,a_{-j},p,p_{-j}\right)$.
where a is the decision to invest in post-Panamax vessels, F_j is the total cost of port j increasing capacity to post-Panamax, σ is the fraction subsidied by the Army Corps of Engineers or some other outside authority, and a_{-j} is the vector of investment decisions by other ports. The decision to invest is affected by the expected difference in W but also the amount that is subsidized. Clearly if the entire cost were subsidized, $\sigma_j=1$, ports would always invest.
This model is static; the authorities do not consider their decisions in the future. However, the status of being Big Ship Ready is persistent. For ports which are already Big Ship Ready entering the period, $\omega_j=1$, I assume that $F_j=0$, and so the decision is to always "invest" for these port authorities.

1.4.6 Equilibrium

An equilibrium is a vector of carrier prices, quantities, ship sizes, and investment decisions, $\{\left(p_{sea,o,j}\right)_{o,j},\left(p_{land,j,d}\right)_{j,d},\left(q_{o,j}\right)_{o,j},\left(S_{o,j}\right)_{o,j},\vec{a}\}$ such that

1. Taking prices and investment decisions as given, importers solve equation (1.1).

2. Taking importer quantities and investment decisions as given, carriers solve equation (1.2).

3. Taking ship sizes, demand, subsidies, and the investment decisions of other ports as given, ports solve equation (1.4).

1.5 Estimation

There are three sets of parameters of interest: the economies of scale, the substitution parameters, and the port authority weight on total quantity. Identifying substitution parameters requires costs that depend on economies of scale, and identifying how much authorities weigh quantity requires understanding demand. I present how I estimate each set in this order.

1.5.1 Economies of Scale

Carriers that use larger vessels have a lower average cost per container, and they pass these savings onto importers. I parameterize the average cost function as

$$f\left(S_{o,j}, D\left(o,j\right); \beta_{sea}\right) = \beta_{sea} \times S_{o,j}^{-\gamma} D\left(o,j\right)$$

where γ is the elasticity of average cost with respect to vessel size. Distance $D\left(o,j\right)$ and vessel size $S_{o,j}$ are observed, thus I need to estimate γ and β_{sea}.

If I could see the costs of each ship, estimating γ would be straightforward: simply regress the logged costs on the logged size, perhaps controlling for year built and builder to account for technological progress. Unfortunately, though I see the individual operating costs of ports I do not have those for specific ships. Instead, I use estimates from work I have done with Tom Holmes (Bailey and Holmes (in progress)). In that paper, we use average charter rates published by Drewry to estimate non-fuel operating costs. For fuel costs, we predict the power necessary to move the ship at its design speed based on its physical dimensions. See Appendix for more details.

1.5.2 Demand

I define a market to be all ports in the United States over one calendar year. Although it is not uncommon for carriers to add or remove routes in response to changes in demand, every spring they decide a rough schedule based on anticipated demand and long-term contracts. The port financial reports are issued annually, which makes price vary at the annual level. To make ports comparable, I take averages of values within one calendar year. For example, the Port of New York and New Jersey fiscal year starts January 1, 2015 and the Port of Los Angeles's starts July 1, 2015. For Los Angeles, I treat the 2015 price to be the average of what I gather from the 2014 and 2015 annual reports. Similarly, if a dredging project is completed, say, in October 2016, the 2016 harbor depth is reported as a weighted average of the depth for the first 10 months and the last two.

For simplicity, the price importers pay depends only on distance, port fee, and whether a port is "Big Ship Ready." I define this last characteristic as an indicator for whether a port has a maximum berth depth greater than 45 feet and an air draft above 185 feet. I could interact it with ship size in a binary way, where the ship variable is 1 if post-Panamax and 0 if smaller. Because there is not a strict cutoff, I use a continuous interaction instead. Another possibility would be to estimate this expression and simply remove from the "small" ports from the choice set of larger ships. This would effectively be using a binary variable and taking $\beta_{BSR} \to \infty$. However, even ports that are not technically prepared for post-Panamax ships can take them in under special (expensive) conditions. (See the Data section for an example with the Mediterranean Shipping Company and Port of Savannah.) Thus including all ports in the choice set of all vessels and allowing β_{BSR} to be a large but finite value is a better approximation than imposing a sharp cutoff.

Shipments with ports of entry in cities different than the port of unlading must go through Customs twice. There is likely a fixed cost associated with these ports differing, on top of the additional distance. I include an indicator for ports of unlading and entry differing to account for this cost.

Importer i's utility from going to port j is

$$U_{ij} = -\alpha \left(\beta_{sea} \times \left(S_{o_i,j}^{\frac{-1}{3}} D\left(o_i, j\right) \right) + p_j + \beta_{land} D\left(j, d_i\right) + \beta_{BSR} \mathbb{1}_{BSR} S_{o_i,j} \right) + \xi_j + \varepsilon_{ij}$$

The importer is concerned only with the prices paid to the ocean and land carriers and the "port quality" ξ_j, which is further split into a time-varying component $\tilde{\xi}_{jt}$ and a fixed component $\overline{\xi}_j$. (I drop the t in the rest of the equations to reduce notation and because ξ_j is the only place where it matters.) I do not observe these prices, though I do observe cost components like distance, ship size, and port amenities. I use variation between the importers and the Panama Canal constraint to identify importer-specific components like distance, and information on port operating expenses to identify the effect of prices.

I follow (Goolsbee and Petrin (2004)) to estimate demand. I use maximum likelihood to estimate the parameters for importer-level characteristics, like land and ocean distance, and then an instrumental variable regression to estimate the parameters for port-level characteristics, like price and the unobserved quality. First, note that the importer utility can be rewritten as

$$U_{ij} = \delta_j - \alpha \left(\beta_{sea} \times \left(S_{o_i,j}^{\frac{-1}{3}} D\left(o_i, j\right) \right) + \beta_{land} D\left(j, d_i\right) + \beta_{BSR} \mathbb{1}_{BSR} S_{o_i,j} \right) + \varepsilon_{ij}$$

where $\delta_j = -\alpha p_j + \xi_j$. The location of importers and where they are importing from is random and exogenous, so there are importers who are the same but for distance to ports, the same but for distance from foreign port to port, etc. This variation allows me to identify $\alpha \times \beta_{sea}$, $\alpha \times \beta_{land}$ $\alpha \times \beta_{BSR}$, and δ_j with maximum likelihood estimation. If there were no unobserved product characteristic common to all consumers, variation from the consumers would be enough to identify the price elasticity. However, there are likely quality difference between ports that I do not capture with just my measures of distance and capacity. For example, some ports may have faster turnaround times, or refueling may be cheaper. Price is likely to be correlated with these unobserved qualities. I instrument price with the average operating expense per container. This gives me α alone, and from this I can separate the other variables estimated with maximum likelihood.

Importer i takes as given the size of the ship going from the foreign port to port j, $S_{o_i,j}$. However, the size is the result of decisions made importer i as well as all of the other importers, and I therefore cannot assume that for a port j' that i decides to *not* choose, $S_{o_i,j'} = S_{o_i,j}$. To account for these differences, I nonparametrically predict[6] the size of the vessel as a function of:

- Big Ship Ready status
- total number of containers moving from a given foreign region to that port (including the container of the importer in question)
- the ocean carrier hired by the importer

and use these predicted values for the sizes the importer faces at each port.
Assuming the idiosyncratic term follows a Type I extreme value distribution, the probability of importer i importing a good from o to d going to port j is

$$\mathbb{P}_{iodj} = \frac{e^{\delta_j - \alpha\left(\beta_{sea}\times\left(S_{o_i,j}^{\frac{-1}{3}} D(o_i,j)\right) + \beta_{land} D(j,d) + \beta_{BSR} \mathbb{1}_{BSR} S_{o_i,j}\right)}}{1 + \sum_{k\in J} e^{\delta_k - \alpha\left(\beta_{sea}\times\left(S_{o_i,k}^{\frac{-1}{3}} D(o_i,k)\right) + \beta_{land} D(k,d_i) + \beta_{BSR} \mathbb{1}_{BSR} S_{o_i,k}\right)}} \tag{1.5}$$

and the share of containers going to j is

$$s_j = \int \mathbb{P}_{iodj}\phi(i,o,d)didodd$$

where $\phi(i,o,d)$ is the probability density of importers of type i importing from o to d.

1.5.3 Port Authority Objectives

Recall that ports maximize a weighted average of profits and quantity,

$$\left[(p_j - mc_j)\, q_j\left(\Omega, a, a_{-j}, p_j, p_{-j}\right)\right] + \theta q_j\left(\Omega, a, a_{-j}, p_j, p_{-j}\right) - \mathbb{1}_{a_j}\left(1-\sigma_j\right) F_j$$

Let $\hat{a}_j$ be the action port j does *not* take. Clearly $W(a_j) - \mathbb{1}_{a_j}(1-\sigma_j) F_j \geq W(\hat{a}_j) - \mathbb{1}_{\hat{a}_j}(1-\sigma_j) F_j$, or, rearranging,

$$\theta_j \geq \frac{(1-\sigma_j) F_j - \left[(p-c)\, q - (\hat{p}-\hat{c})\, \hat{q}\right]}{q - \hat{q}} \tag{1.6}$$

if $a_j = 1$ and $q > \hat{q}$ and

$$\theta_j \leq \frac{(1-\hat{\sigma}_j) \hat{F}_j + \left[(\hat{p}-\hat{c})\, \hat{q} - (p-c)\, q\right]}{\hat{q} - q} \tag{1.7}$$

if $a_j = 0$ and $q < \hat{q}$.

[6] Specifically, I use a random forest algorithm. This technique creates decision trees for many random samples of the data and averages over all of the trees' predictions.

I assume that $\theta = \theta_j$, $\forall j$. Then for the nine ports on the East Coast that were not already Big Ship Ready, I have a system of inequalities that allows me to bound θ.

I assume there is a single θ but allow F_j to vary across ports. (I do assume $F_j = \hat{F}_j$, that is, the "hypothetical" cost is the same as if the cost were realized.) Unfortunately, the only ports I see during this period to finish their Big Ship Ready projects are Miami, Houston, and New York & New Jersey. I use the actual costs for their values of F_j. I also see the projected costs of Savannah, Charleston, and Port Everglades. By the time of writing these projects are closer to completion, so the most recent projected costs are likely the actual ones. For the other four ports on the East Coast I do not observe investment costs and thus do not include them in the estimation. (I do observe physical characteristics and could potentially predict costs, but this would require additional modeling.)

For estimation I allow there to be an additional error term on the right hand sides of equations 1.6 and 1.7. This encompasses measurement error. We may worry there is an unobserved cost component, say ζ_j. In that case a port will *not* invest if

$$
\begin{aligned}
&[(\hat{p}_j - \hat{c}_j)\,\hat{q}_j] + \theta\hat{q}_j - (1-\sigma_j)\,F_j - \zeta_j \leq [(p_j - c_j)\,q_j] + \theta q_j \\
&\Rightarrow \theta + \frac{\zeta_j}{\hat{q}_j - q_j} \leq \frac{(1-\hat{\sigma}_j)\,\hat{F}_j + [(\hat{p}-\hat{c})\,\hat{q} - (p-c)\,q]}{\hat{q}-q}
\end{aligned}
$$

and a parallel expression holds if it does invest. Unfortunately there is not a straightforward way to identify both θ and ζ_j without additional assumptions. One potential assumption would be that ζ_j varies across ports but is constant over time and use the panel aspect of the data. I leave this for future work.

Timing and Dynamics

Though the model abstracts from dynamics, we still need to consider authorities' discounting and expectations to make the payoff W_j (a flow) and cost F_j (a lump sum possibly amortized over several years) comparable. There are a few possible sources for the discount rate. One is the rate the port authorities face in the financial markets: all of these ports sell millions of dollars worth of bonds in liquid bond markets, so we can look at the rates they face there to get a sense of what the an annualized F_j amounts to. These rates vary between around 3-7% across and within ports, with a mode around 5%. There are also standard numbers used by federal agencies in conducting cost-benefit analysis. The Army Corps of Engineers uses a value based on the average yield of long-term Treasury securitities; in 2016 this was 3.125%. Other agencies use different numbers: the Office of Management and Budget has used 7% since 1992 (Army Corps of Engineers (2007)).

For my baseline, I use 5%. This number is the one most tightly linked to prices the individual ports face, and is also in the middle of the 3.125-7% range.
More generally the model is potentially misspecified because it ignores dynamics. For example, I implicitly assume the total number of imports is stationary, but if port authorities expect international trade to grow, I could be measuring as a high value of θ what is really future expected profits.[7] I check the robustness of my results by studying counterfactuals with the measured θ and with $\theta = 0$, i.e. when they are profit-maximizing.[8]

1.6 Data

1.6.1 Port characteristics

Data on port characteristics come from port authority financial reports and various news sources. The physical characteristics I focus on are harbor depth and, if applicable, the air draft or height limit due to bridges. There are other forms of capital ports invest in such as rail connections or warehouse areas, but these two are especially important as they limit the maximum ship size. There is a rarely a hard cap, but something like a shallow harbor depth can make it more costly for a carrier to enter with a larger vessel. For example, in 2006, the carrier MSC started sending 6,700 TEU vessels to the Port of Savannah. The draft of these vessels when fully loaded was 48 feet, more than Savannah's harbor, so MSC was obliged to send vessels that were less than fully loaded. This either meant a more complicated scheduling problem or losing some of the economies of scale advantages of a larger ship. They also required high tide and thus could only enter during specific times of day. Eventually, this proved too cumbersome, and MSC switched to sending many of the larger vessels to Charleston, with goods intended for the Savannah market trucked down (Army Corps of Engineers (2012)).
In Table 1.1 are the range and median ports depths on each coast. During this period, most of the growth was from the East Coast expanding to the same depths as West Coast ports. West Coast harbors are for the most part naturally deeper. There is also no canal constraining ship sizes from East Asia to the West Coast, and so most had already adopted to larger ship sizes by this time.
Table 1.2 shows the level of investment at some of the ports on the East Coast. These investments vary in size, but all involve dredging or raising bridges so that the new ships going through the Panama Canal can berth. I consider a port to be "Big Ship Ready" if

[7] In a 2012 study, the Army Corps of Engineers used predictions from IHS Global Insight of 5.2% per year for the next thirty years.

[8] Although it is unlikely the ports truly expected stationary demand, the data alone suggest it would not be an outrageous assumption. From external sources, the growth rate of total container imports was only 1.5% per year from 2007 to 2017 (versus 0% implicity assumed here).

	East Coast		West Coast	
	Max	Median	Max	Median
2014	50	43	53	51
2015	50	44	53	51
2016	51	45	53	51
2017	55	45	53	51
2018	55	45	53	51

Table 1.1: Maximum Berth Depth (ft.)

	Start Year	End Year	Cost	Port Share
Miami	2013	2015	$220 mil.	49%
Savannah	2015	2022	$973 mil.	25%
NY & NJ	2015	2017	$1.68 bil.	100%
Charleston	2018	2021	$558 mil.	36%
Port Everglades	2020	2022-2025	$389 mil.	49%
Houston	2018	ongoing	$1 bil.	35%

Table 1.2: Selected investments in post-Panamax readiness

it has at least one berth with depths exceeding 45 feet and air drafts below 185 feet. The actual depths of even vessels of the same capacity varies, but 45 feet can accommodate ships with capacities below the pre-expansion Canal and many sizes above it. From 2014-2018, Miami, Houston, and New York & New Jersey all became ready under the definition that their harbor depths exceeded this depth. (Houston was already close to 45 feet and thus became ready early on in the construction.) Norfolk already met this criterion entering the period. However, all ports but Boston and Wilmington, North Carolina began at least the initial phases of expanding and so I consider all but those two to have made the investment decision in the counterfactuals. For ports for which I do not have investment cost data, I assume their cost would be the average of the observed ($690 million) and they would pay one third of this ($230 million).

Prices are another important characteristic. Port authorities are legally required to release general schedules of tariffs outlining every possible fee they may charge carriers. These published tariffs are only upper bounds, though, and in practice ports change them infrequently. As mentioned above, most of ports' container revenues come from wharfage, a per-container fee. (Over 90% of container revenues for Port of Los Angeles were from wharfage.) I therefore derive unit price by dividing container revenue by the total quantity

Figure 1.1: Map of ports in sample

of containers[9]

Perhaps the most important feature of ports is one over which the authorities have no control: where the port is physically located. Figure 1.1 shows the map of the ports that I consider in my market. There are two things to notice. First and most obviously, the North American landmass separates ports on the East and West Coast. Carriers going from East Asia to the East Coast must decide whether to unload on the West Coast and ship by land or go through the Canal, subject to the constraints that imposes. Second, ports are not distributed uniformly. The East Coast is more densely filled than the West. There are also many “twins,” ports that are separate but very close to one another, like Los Angeles and Long Beach, Seattle and Tacoma, or Miami and Port Everglades. These facts suggest that substitution patterns likely vary across ports.

1.6.2 Container imports

All imports to the United States file a bill of ladings with United States Customs and Border Protection (CBP).[10] These contain the date entering the U.S., the vessel the shipment

[9] The total quantity here refers to imports *and* exports. I do not study exports in the rest of the chapter but ignoring them here would cause me to overshoot the average price charged.

[10] I thank Tom Holmes for providing these bill of ladings data.

came on, the foreign port it came from, the U.S. port it entered, and the U.S. customs office it went through. These last two variables, the "port of unlading" and "port of entry," respectively, are important in proxying for the land distance traveled. Each observation is at the level of a shipment, which normally consists of all the goods of a certain type a specific importer is ordering on a single voyage. Very often a "shipment" and a "container" are the same and in the model and rest of this chapter I use the terms interchangeably.[11]

I merge these bill of ladings data with vessel characteristics and port characteristics. For vessels the most important characteristic is ship size. I use the vessel names and International Maritime Organization (IMO) number (an identifier that never changes over the ship's life) to add the total container capacity taken from `http://www.vesseltracking.net`. I also add the port characteristics described above. Finally I calculate the distance traveled on both legs of the journey, as I describe below.

The purpose of the shipping industry is to turn goods that are far away into goods that are nearby, and so it is critical that the distances I use are realistic. Land distances are the fastest route by car using the Google Maps API. For the sea distances, I use publicly available data on the actual travel distance between ports from `https://sea-distances.org`. For ships that could go through the old Panama Canal, or "Panamax sized ships," I assume they always take the shortest path, which may or may not go through the Panama Canal. For ships larger, called "post-Panamax," if they are also under 13,000 TEU, I assume they take the shortest non-Panama distance before the canal expanded and the shortest distance, including Panama, afterwards. For ships larger than 13,000 TEU, I assume they always take the shortest non-Panama distance.

Tables 1.3 and 1.4 display the average distances traveled by land for containers and sea, respectively. The majority of containers go through customs in the same port as their ship arrives, so the median distance traveled on land is zero. To give a sense of what these distances refer to, 2,211 km is the distance between Los Angeles and Houston. For sea, 10,601 km is the distance between Yantian (a district in Shenzhen in southern China) and Los Angeles. Over this entire period, the distance between the port of unlading and the final destination is falling. The sea distances are growing before 2016, when they fall, but they begin growing in 2018 again. This drop is likely due to containers that might previously have gone on large vessels through Suez instead going through Panama. Over time, the average importer is preferring ports that are closer by land at the expense of sea.

If importers are preferring longer sea travel, it must be because sea travel is becoming cheaper relative to land. In Table 1.5, we see that the size of vessels has been growing over

[11] To account for shipments that are smaller or larger than a single container, in estimation I weigh all observations by the number of containers. This implies that if an importer orders, say, six containers of home goods, the idiosyncratic cost shock ε_{ij} is the same for all six containers.

Year	Mean	Median	Mean (>0)	Median (>0)
2014	462	0	1,994	2,211
2015	414	0	1,936	2,193
2016	391	0	1,894	2,192
2017	371	0	1,816	1,990
2018	367	0	1,795	1,990

Table 1.3: Land distances from port j to destination d (km)

Year	Mean	Median
2014	12,166	10,571
2015	12,459	10,660
2016	12,300	10,601
2017	12,207	10,601
2018	12,334	10,601

Table 1.4: Sea distances from foreign origin o to port j (km)

time. Given the economies of scale in shipping, this leg of the journey has become cheaper.

Year	Mean	Median
2014	6,465	6,350
2015	6,796	6,600
2016	7,061	6,763
2017	7,504	7,500
2018	7,728	8,000

Table 1.5: Vessel size (twenty-foot units)

1.6.3 Summary

Table 1.6 shows the variables used in the estimation of scale economies, demand estimation, and the weight ports put on quantity. I estimate the values in that order, using scale economies parameters as part of the demand estimation, and then using the substitution parameters from demand estimation to simulate market shares under counterfactual scenarios. The differences between simulated and realized profits allow me to bound the weight the authorities put on quantity outside of profit.

	Variable	Source
	Charter rates ($)	Drewry
Scale economies	Fuel costs ($)	West Texas Intermediate
	Port choice	US Customs Bill of Ladings
Demand estimation	Ocean distance (km)	www.sea-distances.org
	Land distance (km)	Google Maps API
	Port fee ($)	$\frac{\text{Container revenues}}{\text{Container quantity}}$ from port annual financial reports
	Ship size (TEU)	www.vesseltracking.net
	Big Ship Ready ($\{0,1\}$)	Port websites, ACE reports, news reports
	Expansion costs ($)	Port websites, ACE reports, news reports
Port objective	Marginal costs ($)	Waterfront Commission of New York
	Port profits ($)	Port annual financial reports

Table 1.6: Variables Used

1.7 Results

1.7.1 Economies of scale

In Table 1.7 I compare my results to past papers.[12] These authors use different data sources: (Jansson and Shneerson (1978)) use data from Zim and the port of Haifa in the early 1960s; (Cullinane and Khanna (1999)) use the Fairplay dataset of vessels from the mid-1990s. Nevertheless, the results are fairly consistent across studies. Cullinane and Khanna's elasticity for fuel is an exception: this may be because they assume design speed increases linearly with the size of the vessel. I use the stated design speed and get results much closer to Jansson and Shneerson, who use actual fuel cost data. My total cost elasticity is 0.28. It is not clear what the "total cost" elasticities from the other papers are, though based on the fuel and non-fuel elasticities of Jansson and Shneerson, theirs is slightly higher than mine. In the work that follows I use $\gamma = 0.28$, but values do not change much moving it from 0.20 to 0.35.

Jansson and Shneerson (1978)	0.28 (fuel), 0.4 (capital), 0.6 (non-fuel)
Cullinane and Khanna (1999)	0.03 (fuel), 0.24 (capital)
Author's estimates	0.23 (fuel), 0.27 (capital), 0.52 (non-fuel), 0.28 (total)

Table 1.7: Average cost elasticities

[12] Although I do not use the "capital" elasticity in my estimation, I include it here to show the results are in line with past work.

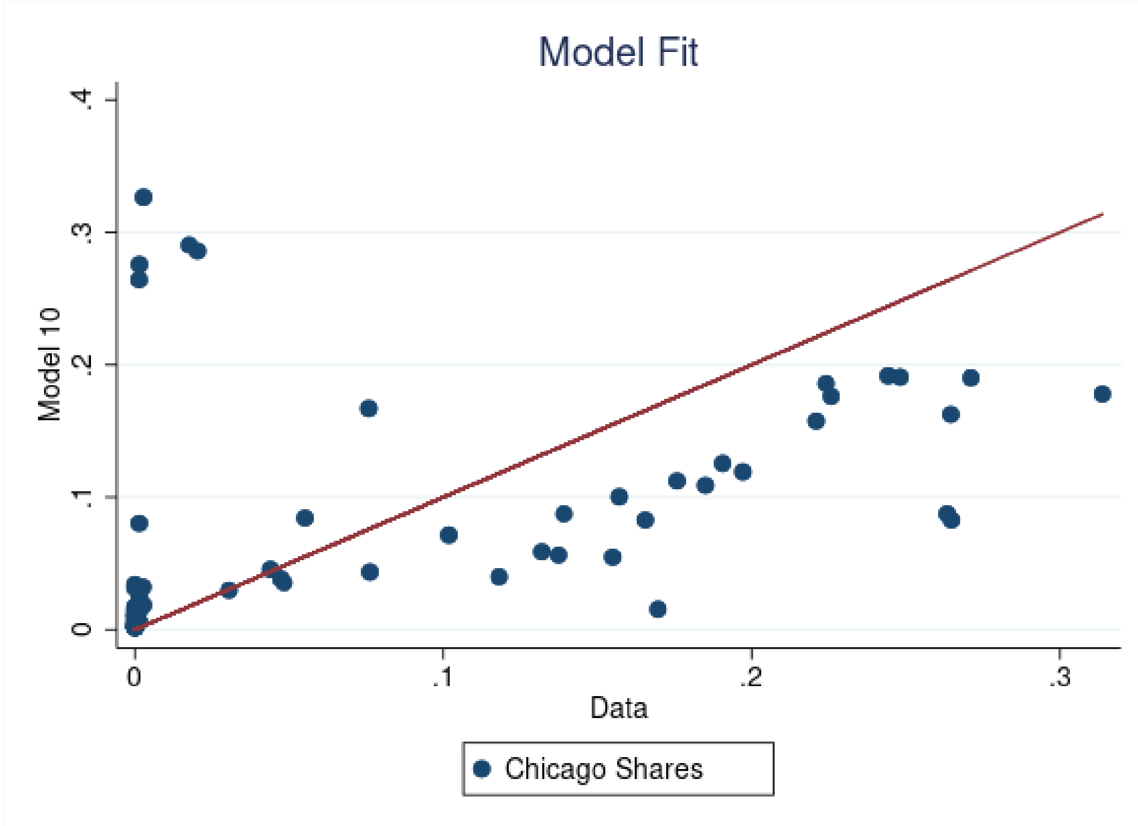

Figure 1.2: Model Fit

1.7.2 Demand

The model fits the data reasonably well. In Figure 1.2, I plot the predicted market share for shipments going to Chicago against the data. I choose Chicago as it is an interior destination geographically closer to the East Coast, but which traditionally sent most of its goods through the West Coast ports. the The model predicts some of the largest ports should be even larger, but is still very close. In Figure 1.3, I aggregate the shares of all East Coast ports and plot that against the actual share. Besides the first year, these shares are also very close.

Table 1.8 shows the demand coefficients estimated. Recall that the coefficients on land distance, sea distance, and whether or not the port can take larger ships show how those variables affect price, not the consumer's utility directly. For that, we would need to multiply them by the price coefficient. E.g., when a consumer is one kilometer farther, his utility increases by $\alpha \times \beta_{land} \times 1 = -0.00116$. In dollar terms we can look at β_{land} alone, so an additional kilometer is about one cent.

It is useful to compare these cost parameters with realistically sized ships. The sea "distance" is scaled down with larger ships to capture economies of scale. Similarly, β_{BSR} is scaled

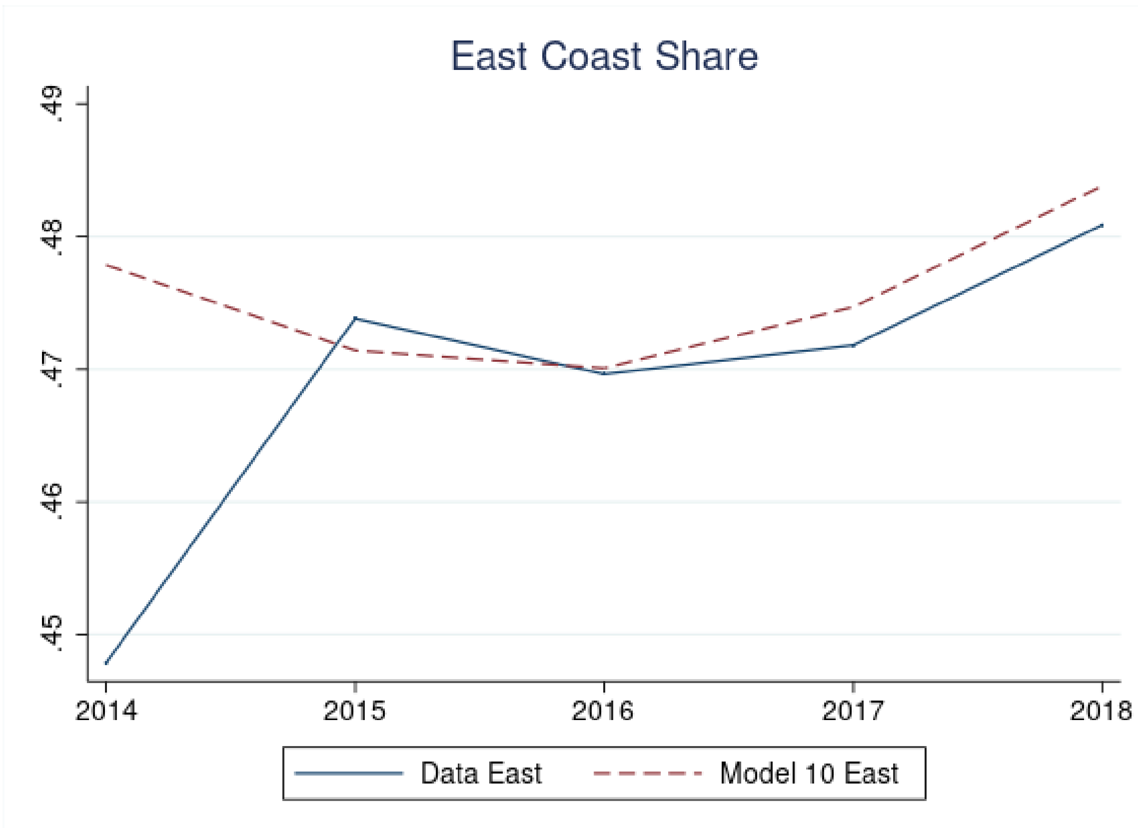

Figure 1.3: Model Fit

up. For a 5,000 TEU ship, distance is multiplied by around 0.09. Thus the cost of shipping by land is about 13 times as expensive as shipping the same container by sea on a 5,000 TEU ship. For larger vessels the difference is even larger. This fits the pattern in the data of importers closer to East Coast ports choosing to use them as the ship size grows, to minimize the total land distance. At the same time, we see that for the same ship, $\beta_{BSR} \times 5000 \approx 176,400$, many orders of magnitude greater than the cost of travel. Why are there not enormous swings in shares when ports becomes Big Ship Ready, then? In Table 1.9 I show some of the time-constant quality effects for each port. These are large in magnitude and they are disparate. Even if a port does not expand, importers persistently value it for other reasons of unobserved quality.

α	-0.0720
	$(*)$
β_{land}	0.111
	$\left(1.14 \times 10^{-5}\right)$
β_{sea}	0.0168
	$\left(2.10 \times 10^{-5}\right)$
β_{BSR}	-4.04
	(0.00460)
$f_{d_i} \neq j$	6.430
	(0.0453)
N	64,020

Table 1.8: Demand parameters

* The objects being regressed in the second stage are themselves estimators. Calculating the standard errors by bootstrapping takes multiple weeks.

In Table 1.10, I present demand elasticities for select ports over this period. I also include the own-elasticity with respect to sea distance. The demand elasticity for sea "distance" also shows what would happen with different sized vessels. For example, plugging in a 5% increase in the size of ships would lead to about a 0.9% increase in Norfolk's 2016 market share. These elasticities are small compared to price and to land. This fits with the findings above that land costs are much higher than sea. Whereas adding or subtracting a few hundred kilometers to the ocean voyage will not have have a large effect on shares, doing the same for land leads to dramatic shifts.

This is further evident in the diversion ratios. Table 1.11 shows the diversion ratios for a sample of ports. These are the ratios if ports in the columns raise their price. For example, if Norfolk raises its price 1%, 33.2% will go to New York and 8.51% will go to Los Angeles.

Port	$\bar{\xi}_j$
25th percentile	-12.224
50th percentile	0.05178
75th percentile	5.660
Los Angeles	-12.197
Long Beach	-12.486
New York	1.0608
Norfolk	-5.399
Houston	4.663

Table 1.9: Port-level fixed effects

Port	Price	Sea Distance	Land Distance
25th percentile	-8.313	-1.262	-3.112
50th percentile	-3.111	-1.103	-2.086
75th percentile	-2.421	-0.928	-1.793
New York, 2016	-3.141	-1.174	-1.828
Norfolk, 2016	-11.969	-1.175	-1.175
Houston, 2016	-8.912	-1.312	-1.964

Table 1.10: Own elasticities for selected ports

It is easy to see how both economies of scale and location affect substitution. Los Angeles and Long Beach are larger than New York, but Boston is so much closer to New York than those two that if it were to raise prices, twice as much would be diverted there rather than the two West Coast ports. Conversely, even though Boston is much nearer Norfolk than Seattle is, there are enough economies of scale from Seattle's size that slightly more traffic is diverted from Norfolk to Seattle than Norfolk to Boston.

1.7.3 Port Authority Objective

Because there are only six moment inequalities, I present them all in Table 1.12 I include the bounds calculated when ports were making decisions in 2017, after the Canal had expanded. I do not observe the investment costs for ports that did not invest, so I can only show lower bounds. The discount rate assumed here is 5%. Using 3.125% or 7% makes very little difference. Miami was the only port to have invested in 2016, but New York and Houston were finished one and two years later, respectively, and the other ports had all announced their plans to expand by this point.

	Boston	New York	Norfolk	Houston	Los Angeles	Long Beach	Seattle
Boston	-	2.10	1.11	0.0611	0.248	0.225	0.229
New York	35.8	-	33.2	3.77	7.44	7.27	6.22
Norfolk	8.36	14.7	-	2.01	2.48	2.36	2.01
Houston	0.214	0.774	0.933	-	3.98	4.17	1.20
Los Angeles	6.43	11.3	8.51	29.5	-	36.2	23.7
Long Beach	6.02	11.4	8.34	31.8	37.3	-	25.8
Seattle	1.10	1.75	1.28	1.65	4.38	4.64	-

Table 1.11: Diversion ratios (%), 2014

Port	Lower Bound
Miami	-105
Houston	290
Port Everglades	106
Savannah	72.5
Charleston	189
NY & NJ	34.1

Table 1.12: Bounds for θ

The lower bounds vary dramatically. The median price-cost margin in 2017 was about $45, so the lower bounds go from a little under this to over 1,400 times as much. The latter value is frankly unbelievable, and is likely a result of the sample being limited to importers in the interior. The fraction of containers with final destination different from the port of unlading is very small for Miami and Port Everglades, 3.4 and 7 percent, respectively, compared to 14.2 percent of containers going to New York & New Jersey or 24 percent for Norfolk. In considering counterfactual revenues and profits I rescale my sample to match the actual total, but ports that have a smaller than average fraction of interior customers will have smaller than actual rescaled revenues. I should emphasize here that the demand estimates are still consistent because the importers are selected on exogenous characteristics. However, there are masses of consumers at each port that likely act differently than interior importers, and it is difficult to say how those affect revenues.

1.8 Counterfactuals

In this section, I imagine there is a centralized East Coast port authority. I first show the extreme cases where there is zero investment and when all the ports invest. The first-order results are not surprising: there would be fewer imports to the East if no one invested and more if everyone did. However, in the world with no investment, East Coast imports fall more than West Coast imports rise. Similarly when everyone on the East invests, imports there rise more than West Coast fall. This suggests that even with the Canal expansion, the two coasts are very imperfect substitutes, and that the East Coast is mostly drawing imports from (or losing them to) the outside good.

More interestingly, I compare the surplus from the actual investment to what might have occurred with coordination. I show the pattern for $\theta = 0$ (port authorities are profit maximizing) and $\theta = 300$, as suggested by the bounds I estimate above.[13] I also study what they would do if they were not subsidized. All scenarios produce results very different from what actually happened. In most, ports working as a single authority would have done better if they had not invested in the Port of New York and New Jersey. In the case where ports are not profit-maximizing and they are not subsidized, they would have chosen to not invest in Houston. These choices align with the ones a social planner considering importer surplus would make.

Predicting counterfactual market shares requires predicting the counterfactual ship sizes, which depend in part on $q_{o,j}$, the quantity of containers going from a foreign port to a domestic one. This of course depends on the counterfactual demand, which is the quantity I'm trying to simulate. I predict ship sizes with a reduce-form model that predicts size based on year, foreign origin, carrier, and importantly, $q_{o,j}$ and whether the port can take in large ships. I estimate the policy functions for p_j in a similar way, as a reduced-form function of investment and demand. These predicted values are used to estimate market shares, which generate new values for $q_{o,j}$. I use these new values to predict new ship sizes, and repeat the process until the $q_{o,j}$ vectors converge.

For the sake of consistency the "Actual" values in the graphs are the model fitted values, not the raw data.

1.8.1 No investment

In 2014, all ports on the West Coast and Norfolk on the East Coast were "Big Ship Ready." By the end of my period in mid-2018, three more on the East Coast had been added: Port of New York and New Jersey, Houston, and Miami.[14] In the first counterfactual, I simulate market shares when the centralized authority decides to have no investments, all ports stay as they were in 2014. The results for the East Coast are shown in Figure

[13] For the investment costs, I use the actual costs for the ports for which I have data. For the ports with missing data, I assume they have costs that are the average of the observed. Results do not differ if I assume they have the lowest costs on record; though there are some changes if I assume they all have the highest costs on record.

[14] Not included in my sample are the Port of Philadelphia, which dredged its main shipping channel to 45 feet in 2017, and Port of Baltimore, which was already 50 feet in 2014. Each of these had 1.80% market share in 2018.

1.4. Without investment, the major East Coast ports are unable to take advantage of the expanded canal. The West Coast share in Figure 1.5 is higher, but less than the decrease in the East Coast.

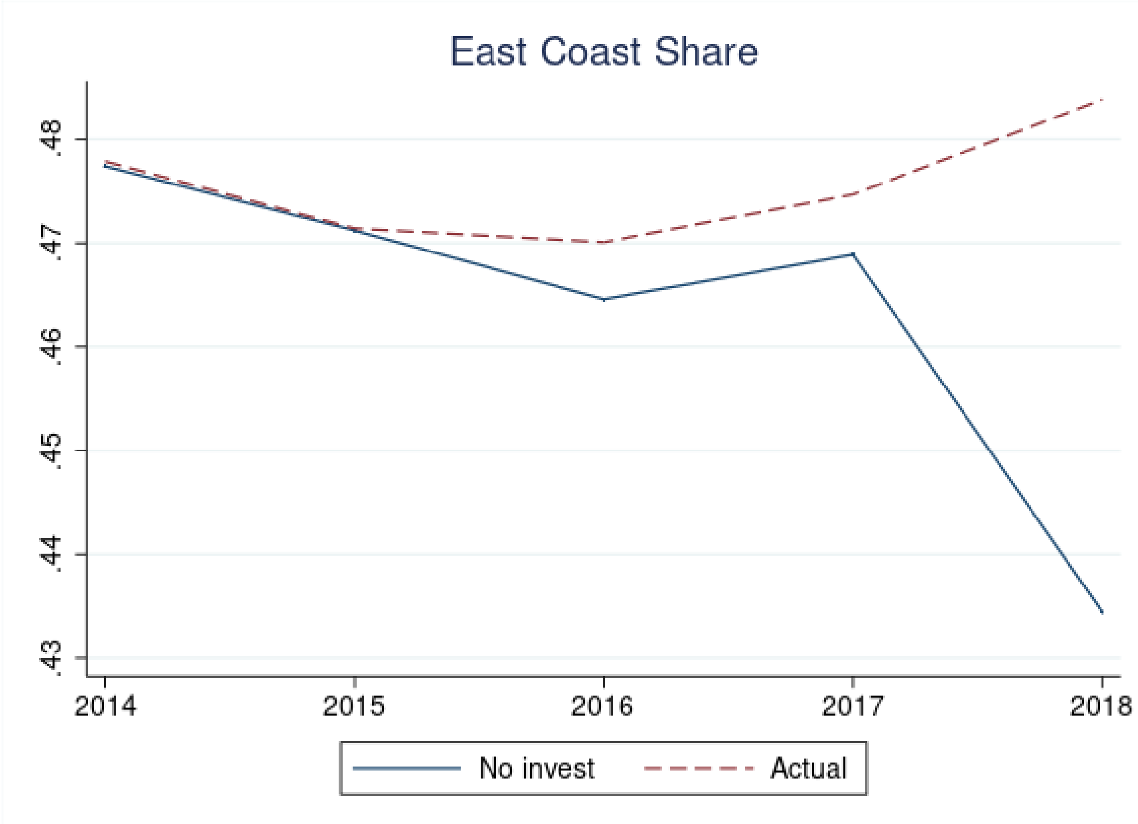

Figure 1.4: East Coast shares when no East Coast ports invest

1.8.2 All invest

In this counterfactual, I consider the outcome where all ports on the East Coast become ready for larger vessels by 2017 at the latest. (Ports that were ready by 2017 are still ready earlier.) Figures 1.6 and 1.7 show the effects in the East and West, respectively. There is a large increase in 2017, though by 2018 the East Coast share is only about 6% higher than it would be otherwise. The West Coast 2018 share difference is smaller, about 4% below its actual value. Just as in the scenario with less investment, the traffic lost does not all go to the West Coast, in this scenario where there's more investment, not all the traffic gained comes at the expense of the West.

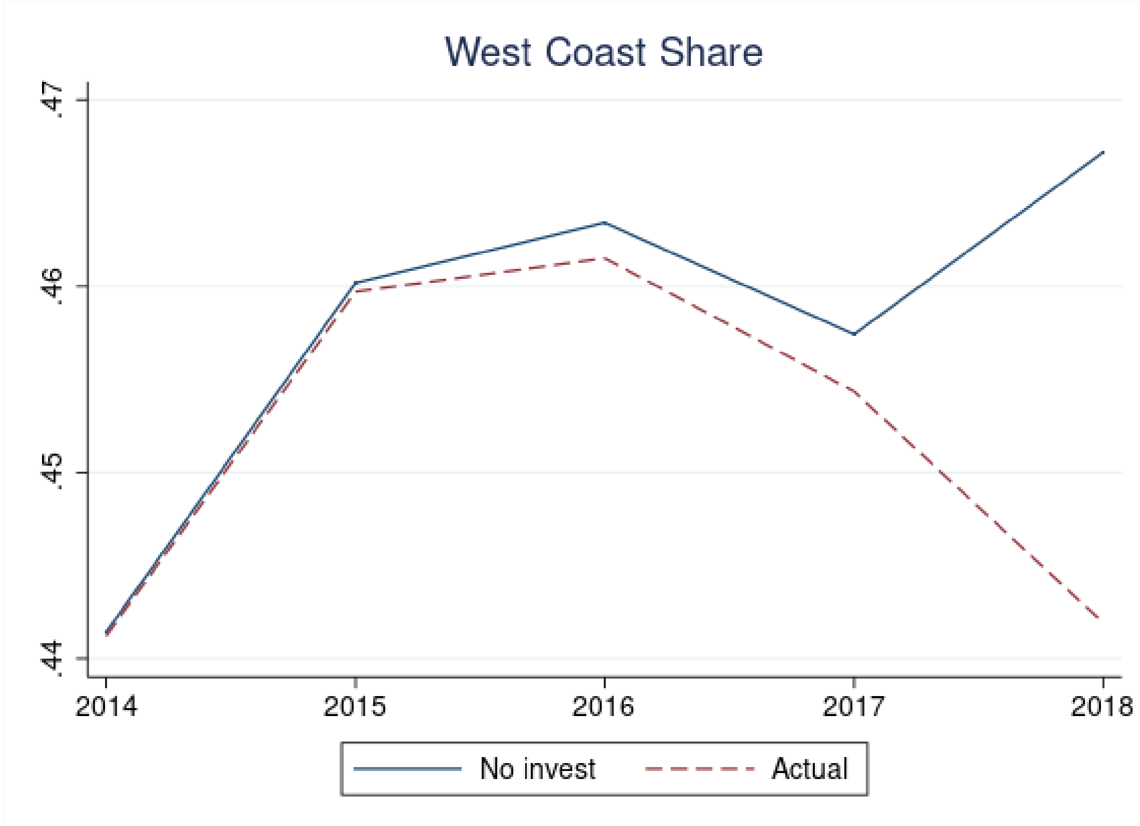

Figure 1.5: West Coast shares when no East Coast ports invest

1.8.3 Investment under a single authority

Suppose now that the East Coast ports wish to choose the investment that maximizes their total surplus. From my estimated bounds, 300 is a plausible value for θ, and so for this counterfactual I assume that port authorities value each container as worth \$300, on top of their unit profit. To show the importance of considering the non-profit maximizing motives of the authorities, I also show results from $\theta = 0$, that is, when port authorities act as normal profit-maximizing firms. The results are in Table 1.13. The table shows port, importer, and total surplus, as well as the port that the centralized authority would *not* invest in (not counting Boston and Wilmington). The table does not show the welfare of the outside government providing the subsidy.

If ports are subsidized and value containers separately from profit, total surplus is \$1.2 billion higher when they do not invest in the Port of New York and New Jersey. If they are solely profit maximizing the difference is even greater, \$1.5 billion. The importers' consumer surplus is about \$22 million when all but Boston and Wilmington invest and is increasing with greater investment, so the savings to the ports alone is even higher. For comparison, in 2017 the ports in the East Coast sample earned about \$1.3 billion in revenue

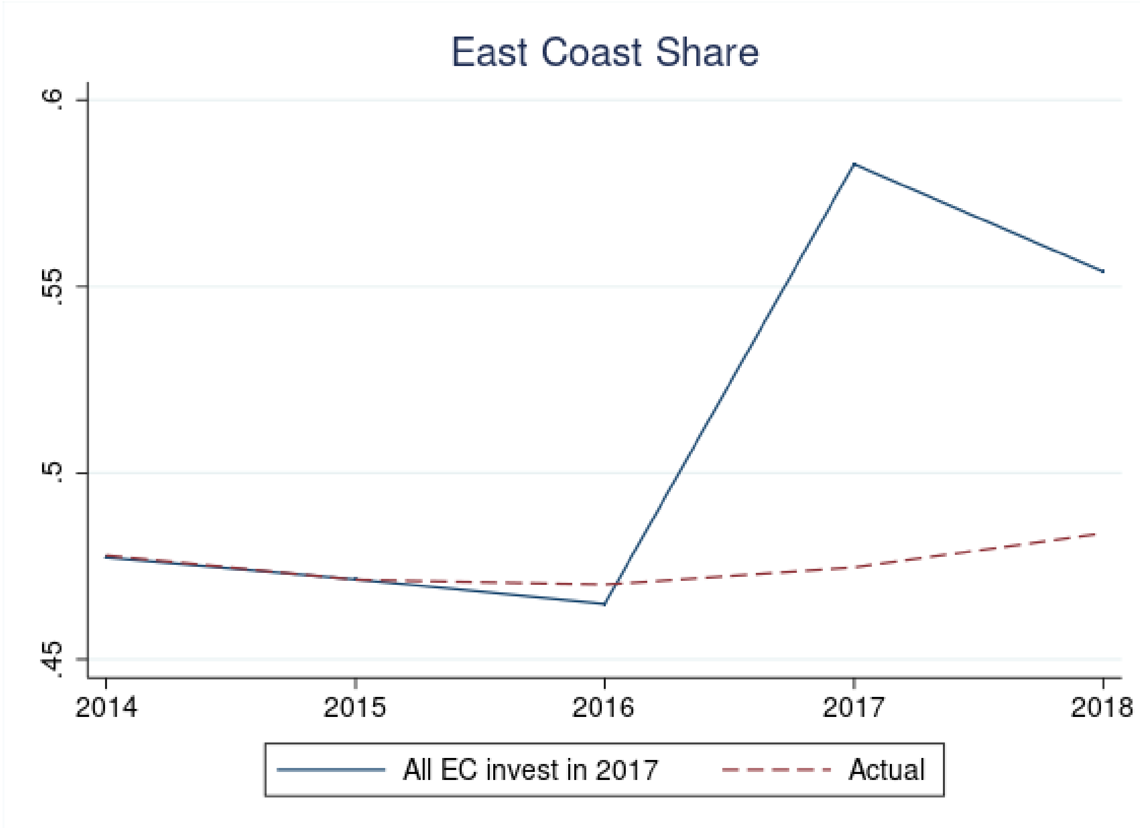

Figure 1.6: East Coast shares when all East Coast ports invest

from importing 11 million containers (plus exporting a slightly lower number). The social welfare loss is therefore almost the equivalent of one year's total revenue. It is important to note these investment patterns are not necessarily the true optimum. Finding that would require resolving the equilibrium for each of the 512 possible combinations, which is computationally slow. However, even considering slight deviations from what ports actually did, we see that there are large savings.

It may be surprising that the ports dropped from investing are ones as large as New York or Houston. There are two forces pushing towards these as the ones to not bother expanding. First, the investment costs for these ports were higher than the others, and in the case of New York, all of this cost was borne by the port authority itself.[15] Second, these ports were already desirable to importers for other reasons. Looking at Table 1.9, both have better than median quality fixed effects. "Port quality" and expansion act as substitutes, to some degree.

Subsidization mechanically affects the total cost of the ports, but it only affects the decision

[15] New York has received federal assistance in the past for dredging projects, but the Bayonne Bridge raising was self-funded.

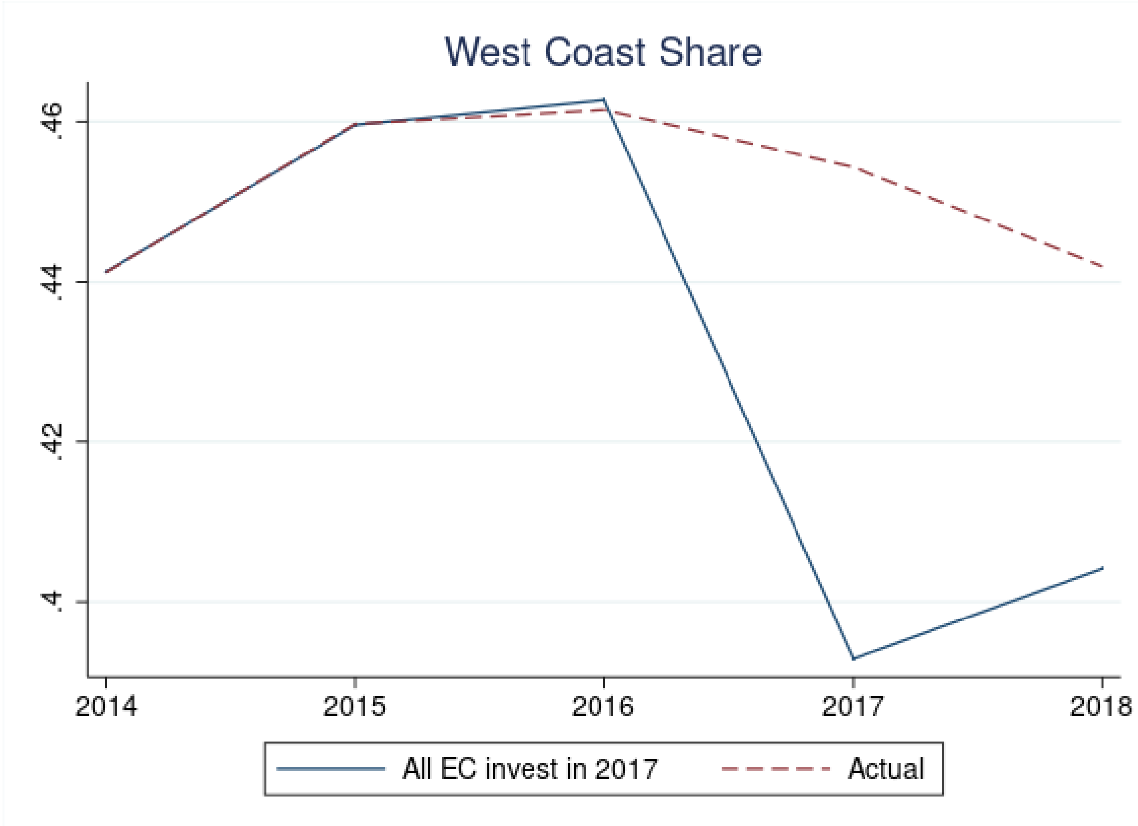

Figure 1.7: West Coast shares when all East Coast ports invest

in the case when ports are profit-maximizing. This makes sense if we consider why a subsidy may exist in the first place: the federal and state governments believe there are benefits from improved trade infrastructure that the port authority cannot charge for, and thus they want to encourage these positive spillovers. However, θ greater than zero implies the authorities are already aware of some of these spillovers and value them in their decision making. Even if the authorities cannot capture the externality through increased revenue, they value it for its own sake.

Finally, in Table 1.14 I show the individual payoffs for New York under the realized outcome (when almost all ports invested) compared to the "optimal[16] " scenarios where New York does not. In the baseline scenario, when $\theta = 300$, New York does better when it chooses to invest if everyone else is, \$397 million versus \$165 million. Even though the raising the bridge was incredibly expensive, over \$1.7 billion, it is worth it to New York if almost everyone else is investing. Thus, there is a "race to the bottom" where individual ports find

[16] I put "optimal" in scare quotes because, though it does happen that this single port authority solution is optimal, there is no theoretical reason why that need to be the case. It is possible that under a different specification, say if the investment decision were continuous rather than binary, the single authority solution would not maximize total surplus.

	Subsidized		Unsubsidized	
	$\theta = 300$	$\theta = 0$	$\theta = 300$	$\theta = 0$
Not investing in...	NY & NJ	NY & NJ	NY & NJ	Houston
Total port welfare ($ mil.)	1,481	-774	-2,649	-5,066
East Coast CS ($ mil.)	19	19	19	21
Total surplus ($ mil.)	1,500	-755	-2,630	-5,045
Actual surplus ($ mil.)	306	-2,239	-3,387	-5,932

Table 1.13: Investment Decisions Under East Coast Authority

		Subsidized	
		$\theta = 300$	$\theta = 0$
NY & NJ	Doesn't invest	165	20
	Invests	397	-789

Table 1.14: Individual Payoffs of New York under "Optimal" Equilibria ($ mil.)

it in their interest to overinvest, even as total surplus falls.

Robustness

Whether or not a port (or coalition of ports) has an incentive to invest is clearly highly dependent on the value of θ. The selection of 300 is reasonable given the governance of the ports and the data observed, but it remains true that the sample size for choosing it is very small. In Figure 1.8, I plot the "best choices" for the Port of New York and New Jersey, the collaborating East Coast, and the collaborating East Coast without any state or federal subsidies. A value of zero means the entity should not invest, one means it should. In the extreme, a very high value of θ implies as many people as possible should invest. In the graph we see that it is in the best interest of New York and also the East Coast coalition for it to invest if θ is over $870. Similarly, if θ is very low, it is in neither's interest to invest. In between, though, from about $210 to $870, is an intermediate range where the incentives are misaligned. My chosen value of 300 falls in this range, but so do many others, and the fundamental result stays the same.

Figure 1.8 also partially answers a question I do not directly address in this chapter: how much of the overinvestment is due to competition and how much is due to the subsidies. The individual best choice for New York is far from the East Coast solution with subsidies, but much closer when there are none. They are still different, to be clear: these best response functions are not probabilistic objects and so there is not chance they are the same line. If we think of distance to the optimal allocation as the level of taxes or subsidies necessary to adjust individual behavior, the tax on investment for New York is much smaller when there are not subsidies than when there are.

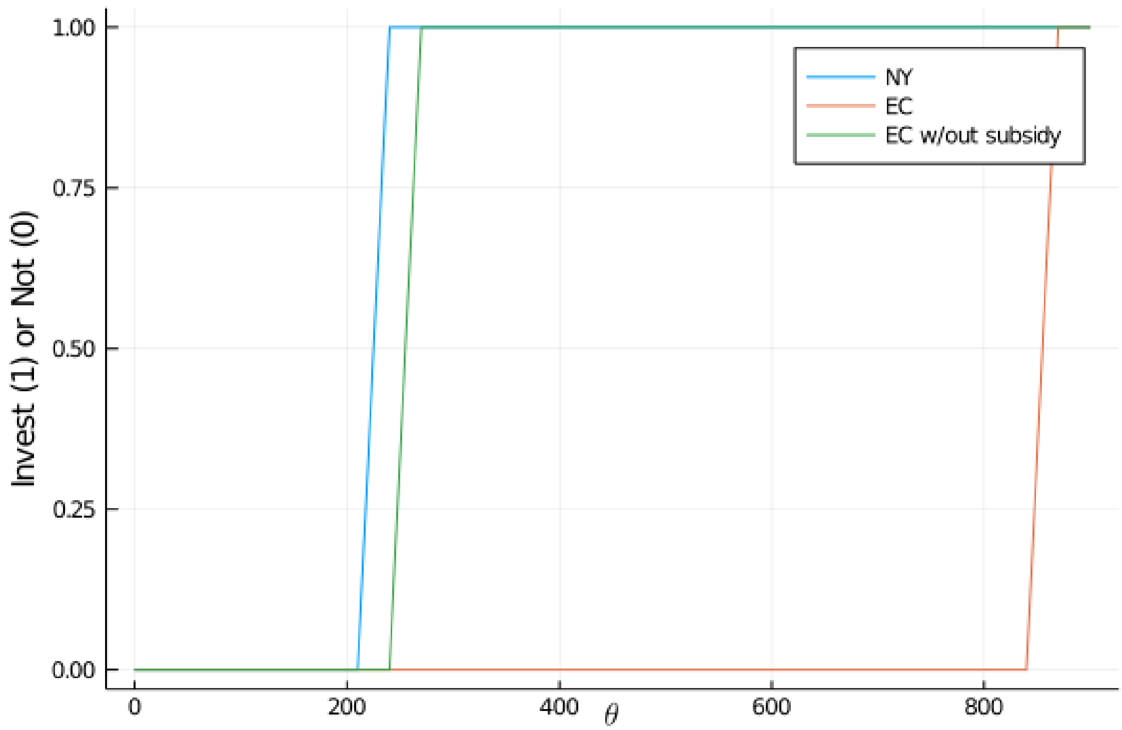

Figure 1.8: Incentives for NY & NJ, all East Coast, and all East Coast w/out subsidy to invest

1.9 Conclusion

Most infrastructure investment is not done by the US national government but smaller political units. However, the effects of this investment are often national in scope. Though competition can have many beneficial effects, such as increased variety, if these units act strategically there may be negative externalities that could be better internalized by a centralized authority.

I have looked at the case of container ports in the US in the period leading up to and right after the Panama Canal expanding. The expansion pushed many ports on the East to invest in larger harbors so they could take full advantage of the larger ships. If their business expanded only at the expense of already large ports, the social cost of the investments would not necessarily exceed the benefits. As it turns out, most of the new business came not from the already large West Coast ports but the smaller fringe, ports that show up as my outside good. The total level of investment may be close to optimal. However, on the East Coast there are clear benefits of coordination, both for the port authorities and society at large. If the port authorities were to coordinate their decisions, they would produce results better for them and better for total social surplus by avoiding expensive investments that do not sufficiently raise the coastwide market share. Subsidies could potentially solve this issue, but as they are currently used appear to be too indiscriminate. If the Army Corps of Engineers plans to continue subsidizing under the same guidelines it currently uses, encouraging ports to coordinate could be one way to produce savings for both the ports and the federal government without too greatly decreasing surplus for importers.

Chapter 2

Work Rules and Port Productivity

2.1 Introduction

Union supporters and opponents alike agree their effects go far beyond wages. The work rules governing what workers and employers can and cannot do are often just as contentious as the pay itself, if not more. Freeman and Medoff (1984) famously argued this could benefit productivity, as giving workers a voice offers an alternative to leaving a disappointing firm, preventing costly turnover. However, there is also evidence that these rules may be an inefficient way of splitting rents and their removal can lead to lower prices and more productive firms Schmitz (2005).

This chapter studies a very particular change in work rules and how the industry's productivity was affected. In 2008, the International Longshore and Warehouse Union (ILWU) agreed to a contract that allowed all ports on the West Coast of the United States to fully automate, even knowing that automated terminals would lead to lower labor demand. Given the union's historical behavior, there is good reason to think employers would not have been able to automate without this provision. Though I am studying a particular contract in a particular industry, the question the provision addresses–what kinds of machinery a firm can invest in, and what happens to workers it may displace–exists across almost all industries and has been a cause of dispute since at least the Industrial Revolution.

I measure the impact in two ways: first, I do a simple event study comparing the average number of worker-hours necessary to unload a ship after the contract was signed. It takes about 25% less worker-hours to unload a vessel after the new contract than before, with the difference coming roughly equally from fewer workers and from fewer hours spent working. This difference decreases after the 2015 contract, though productivity remains higher than it was before 2008. This is a little surprising just looking at the words of the 2015 contract, but I discuss how the change may reflect a more aggressive attitude taken by the ILWU, as signaled by the slowdown leading up to the contract.

A weakness of the previous approach is there may of course be a time trend unrelated to the contract. Including a linear time trend mitigates the issue but does not eliminate it. Because the union contract (by construction) covers the entire West Coast, I cannot find any nearby ports unaffected. I can, however, study what happens to the port of Long Beach when it actually built an automated terminal in April 2016. I estimate a difference-in-difference and find that the differential impact of Long Beach having an automated terminal was about 15% less worker-hours. Before automation, Long Beach actually used slightly more labor than average. After the new terminal was built it used less than average. It is worth noting that the differential effect of automation need not be the same as the coast-wide effect of the new contract, but the fact that both are qualitatively and quantitatively similar is reassuring.

The union contract was not the only change during this period: the average ship in my sample grew from a little under 45,000 gross tons to over 60,000, over 33% growth. Larger ships can carry more goods but also take longer to unload. The additional cargo and additional time do not necessarily grow at the same rate. If they do not, then the aggregate labor productivity (coast-wide containers over coast-wide worker-hours) will be biased because it fails to account for changes in inputs. In my data I observe work times for individual ship dockings. Though results with the aggregate and granular data are qualitatively the same direction, I show that having this fine data is necessary for quantitatively measuring productivity over time.

This chapter follows two major strands of the literature. The first studies how unions and work rules affect productivity. Two papers in particular are closely related to mine. Benjamin Bridgman studies how longshore unions limited productivity growth from containerization Bridgman (2018). Even though labor productivity increased, total labor costs fell by much less until much later on. Rather than study how a union reacts to technological change, I am measuring the effects of a change in the work rules themselves. In his study of US and Canadian iron ore mines, James Schmitz finds productivity increased and prices fell Schmitz (2005). He concludes that the threat of competition forced firms and unions to redo work rules that had been sharing profits at the cost of efficiency. There are some interesting similarities between West Coast ports and US and Canadian iron ore mining. Negotiations for the previous contract had been contentious, leading up to a ten day lockout in 2002. Ocean carriers and other logistics firms started to explore alternatives to West Coast ports. Additionally, in 2006 the Panamanian government had announced it would expand the Panama Canal by 2014[1], making the East Coast ports a better substitute.

The second strand studies how measured productivity may differ depending on the level of

[1] The actual end date was in July 2016.

aggregation. Basu and Fernald (2002) outline a general theory of the differences between measured aggregate productivity and pure technological change. An early empirical paper in this literature, Griliches (1963) studies how changes in input quality and economies of scale may bias results. The studies usually emphasize how firm-level production parameters need not be identical to aggregate parameters because of between-firm reallocation. I am saying little about interactions between firms (or ports, in this case) and focusing more on issues of aggregation within a firm. In particular, how changes in the quality of what is essentially an intermediate product, ships, can affect the measured productivity if not accounted for.

2.2 Background

Since it began in the 1930s, the International Longshore and Warehouse Union (ILWU) has been one of the most powerful labor unions in the United States. It has been able to guarantee high wages for its workers (average full-time salaries are over $180,000), but in this chapter I focus on its jurisdiction over what investments are allowed in its work spaces. Part of the union's success comes from its measured approach to new technology. When containerization began in earnest in the 1960s, the union did not oppose the changes, but bargained such that existing members would continue to receive benefits and pensions even as the demand for their labor decreased. (The fact that port labor costs did not fall is the main reason Bridgman argues containerization did not lower trade costs until years after it began.) It has continued that tradition more recently. In the 2008 contract, in exchange for $800 million in wages and pensions, the ILWU agreed to include the following paragraph:

> **1.72** It is recognized that the introduction of new technologies, including fully mechanized and robotic-operated marine terminals, necessarily displaces traditional longshore work and workers, including the operating, maintenance and repair, and associated cleaning of stevedore cargo handling equipment. The parties recognize robotics and other technologies will replace a certain number of equipment operators and other traditional longshore classifications. It is agreed that the jurisdiction of the ILWU shall apply to the maintenance and repair of all present and forthcoming stevedore cargo handling equipment in accordance with Sections 1.7 and 1.71 and shall constitute the functional equivalent of such traditional ILWU work. [...]

For the first time, the contract explicitly allowed full automation on terminals. Not only was automation allowed in principle, but the union acknowledged that automation could go on even if it was known to cost union jobs.

A written contract is one thing, whether it makes a difference in behavior is another. Port terminals may have had no intention to ever automate terminals. Conversely, the ILWU may have never had any intention of opposing automation if it had not been explicitly allowed. These possibilities seem unlikely. The first is the most obviously false: port terminals did automate. They used technology that was known about in 2008 and that other ports had already implemented. The second, that even without this provision, the union would not have opposed automation, also seems wrong for a similar reason. Even with the provision in place, the ILWU has vocally opposed automation has it has happened. In 2020, ILWU president Willie Adams urged California governor Gavin Newsom oppose automation there Tirschwell (2020) . In Washington State, the union has fought against state grants for replacing existing machine with low-emission ones that also use less labor Editorial (2021). What the union has *not* done is stage strikes or otherwise disrupt work on the docks, which would be a violation of the terms. There is therefore good reason to assume employers and employees take the contract as binding.

The contract for the period at the end of the data is not the 2008 contract; it was one signed in 2015. The 2015 contract did not make as many sweeping changes and it kept the paragraph on automation. It is important to consider, though, because even if the contract itself was unremarkable, the negotiations were far more aggressive than in 2008, culminating in a work slowdown in early 2015. Figure 2.1 shows the total number of vessels coming to the West Coast. The effects of the slowdown are visible in the 200 fewer vessels appearing than there otherwise would have been. When I'm studying the effects of these contracts the written clauses are important, but so, too, are the signals each party may give in how it plans to operate for the coming years.

2.3 Model

I build a simple production function as a framework for understanding productivity, changes to it, and how aggregation could affect its measurement.

2.3.1 Production function

Ports convert the inputs of labor, port infrastructure, and ships into the output of imported goods. The total amount of goods a ship unloads is a function of the size of the ship and the port where it is unloading. For unloading i at time t,

$$Q_{it} = B_{pt} S_{it}^{\eta} e^{\varepsilon_{size,it}} \tag{2.1}$$

where Q_{it} is the quantity to be unloaded, S_{it} is the size of the ship, and B_{pt} is a port p,

time t fixed effect. $\varepsilon_{size,it}$ is an i.i.d. demand shock or measurement error.
For a ship unloading Q, the terminal operator chooses the labor necessary to unload the ship. Port infrastructure and ship size are taken as given by the operator. The operator chooses the minimum L_{it} necessary to solve

$$Q_{it} = L_{it}^{\alpha} K_{it}^{\sigma} S_{it}^{\gamma} e^{\omega_{pt}+\nu_{it}} \tag{2.2}$$

where L_{it} is labor, K_{it} is capital, S_{it} is the size of the ship, ω_{pt} is a port-level productivity term, and ν_{it} is a random productivity shock (unobserved by everyone). Note that the ship size appears here as well as equation (2.1). The rationale is that ship sizes determine not only how much exists to be unloaded, but also the rate of unloading. For example, if there are fixed costs of unloading any ship, a larger one may be unloaded faster than two smaller vessels carrying the same quantity. There may also be diseconomies. Equations (2.1) and (2.2) are flexible enough to allow both.
L_{it} is clearly a function of the other variables, and because Q_{it} is a direct function of size[2] , we can rewrite it as

$$L_{it} = \Big(\frac{B_{pt}}{e^{\omega_{pt}+\nu_{it}}} S_{it}^{\eta-\gamma} K_{it}^{-\sigma}\Big)^{\frac{1}{\alpha}}$$

Logged, this becomes

$$\begin{aligned} \ell_{it} &= \tilde{b}_{pt} + \frac{(\eta-\gamma)}{\alpha} s_{it} - \frac{\sigma}{\alpha} k_{it} + \tilde{\omega}_{pt} + \tilde{\nu}_{it} \\ &\equiv \frac{(\eta-\gamma)}{\alpha} s_{it} + f_{pt} + \tilde{\nu}_{it} \end{aligned} \tag{2.3}$$

where lowercase letters signify logs and $\tilde{x} = \frac{1}{\alpha}x$. (Because $\varepsilon_{size,it}$ and ν_{it} are both i.i.d. I absorb the former in the latter.) I further assume that all ships entering a given port in a given month have the same access to capital, so $k_{it} = k_{pt}$ and I collapse all the port-time level variables into a single variable f_{pt}.
Logged labor productivity is

$$q_{it} - \ell_{it} = \frac{(\alpha-1)\eta+\gamma}{\alpha} s_{it} + g_{pt} - \tilde{\nu}_{it} \tag{2.4}$$

where g_{pt} is a port-time level effect. Unless $(\alpha-1)\eta+\gamma = 0$, there are thus scale effects from the size of the ship. If $(1-\alpha)\eta < \gamma$, there are economies of scale, if $(1-\alpha)\eta > \gamma$, there are diseconomies.

[2] There is some abuse of notation here. The relationship in equation (2.1) relates the stock of quantity to the stock of size, whereas the production function also considers time (i.e., total worker-*hours* rather than total workers.

2.3.2 Effects of new contract

The new contract affects production through a combination of new capital (k_{pt}) and port-wide productivity (ω_{pt}). These changes would show up as level shifts through the g_{pt} term in equation (2.4). For most of the series, I cannot distinguish between demand shifts ($\tilde{b}_{pt}$), capital investment, and productivity changes, so changes to the combination of all of them, g_{pt}, is the object of interest. After a new contract in year, the port level productivity g_{pt} becomes $g_{pt} + c_y$. This c_y is the effect of the contract, not necessarily automation itself. In the case where I do observe a terminal becoming fully automated, I allow that to have a different effect.

2.3.3 Aggregate versus average productivity

"Aggregate" productivity is defined as the ratio of coast-wide output over coast-wide inputs. "Average" productivity is the average of the ratios of unloading-level output over unloading-level inputs. Neither measure is more "correct" than the other but their usefulness depends on the question being asked. (See Basu and Fernald (2002) for examples of how aggregate may be useful even when it does not reflect technological change.) I am interested in a change in technology, broadly defined, rather than between-port reallocation or changes in inputs, and therefore need to measure average. Because of scale economies and changes in inputs, these values may not be the same.

A simple decomposition shows the relationships between the two. For each unloading i, Q_i is the quantity unloaded and L_i is the labor employed. The labor productivity of unloading i is simply $A_i \equiv \frac{Q_i}{L_i}$. Let $\mathcal{A}$ be the aggregate labor productivity, i.e.,

$$\mathcal{A} \equiv \frac{Q}{L} \equiv \frac{\sum_i Q_i}{\sum_i L_i}$$

It is clear that aggregate productivity is not the simple average of ship-level productivities. Rearranging, we get

$$\begin{aligned} \frac{\sum_i Q_i}{\sum_i L_i} = \frac{\sum_i \frac{Q_i}{L_i} \times L_i}{\sum_i L_i} = \frac{\sum_i A_i \times L_i}{\sum_i L_i} = \frac{\mathbb{E}\left[A_i L_i\right]}{\mathbb{E}\left[L_i\right]} = \frac{Cov\left(A_i, L_i\right)}{\mathbb{E}\left[L_i\right]} + \mathbb{E}\left[A_i\right] \\ \Rightarrow \mathcal{A} = \mathbb{E}\left[A_i\right]\Big(corr\left(A_i, L_i\right) + 1\Big) \end{aligned}$$

If the unloadings with higher labor productivity A are correlated with more labor employed, the measured aggregate productivity $\mathcal{A}$ will be higher than the average A. If higher labor productivity is correlated with less labor, the opposite will be true. Returning to the parameterized production function: holding all else equal, if there are decreasing returns for labor ($\alpha < 1$) we would expect to see unloadings with more labor to have lower productivity:

Source	Variables
PMA dispatch summaries	vessel hours at port, total hours worked per vessel, total gangs employed per vessel
PMA annual reports	aggregate hours, aggregate tonnage
www.vesseltracking.net	vessel sizes
Bills of lading	containers unloaded per vessel

Table 2.1: Data Sources

$corr\,(A_i, L_i)$ is negative, and $\mathcal{A} < \mathbb{E}\,[A_i]$. If there are increasing returns, we would expect to see the opposite, $\mathcal{A} > \mathbb{E}\,[A_i]$.

I am interested in changes in productivity. If there were only a level difference between $\mathcal{A}$ and $\mathbb{E}\,[A_i]$ that was constant over time, looking at $\mathcal{A}$ would be sufficient. However, labor is not the only input. Ships sizes are changing during this period, which means $corr\,(A_i, L_i)$ is also changing over time. To accurately measure the average productivity, I need unloading-level data.

2.4 Data

Table 2.1 lists the data sources I use. The dispatch summaries from the Pacific Maritime Association (PMA) are the most important as these contain fifteen years (2005-2020) of labor data for four ports on the West Coast of the United States: Seattle, Tacoma, Oakland, and Los Angeles/Long Beach. (Los Angeles and Long Beach are actually two separate ports, but they are close enough to draw on the same pool of labor.) I match the dispatch summaries with data from www.vesseltracking.net for vessel sizes. I also match the vessel sizes to bill of ladings from 2007-2015, in order to measure the relationship between ship size and quantity unloaded.[3] I would ideally match the dispatch summaries with the bills of lading, but differences in how the two datasets record when a vessel enters a port makes the matching non-trivial.

In Table 2.2, I list summary statistics for some of the main variables. My main measure of interest will typically be containers per gang-hour conditioned on vessel size. I do not observe the relationship directly, but I am able to observe the relationship between containers and size and between gang-hours and size, and from these I can indirectly measure the relationship between containers and gang-hours.

Not visible in the table but important for my story are time trends taking place during

[3] Bills of lading are more or less receipts for transporting goods by ship. I use data originally prepared for Holmes and Singer (2018) and thank Tom Holmes and Ethan Singer for sharing their data.

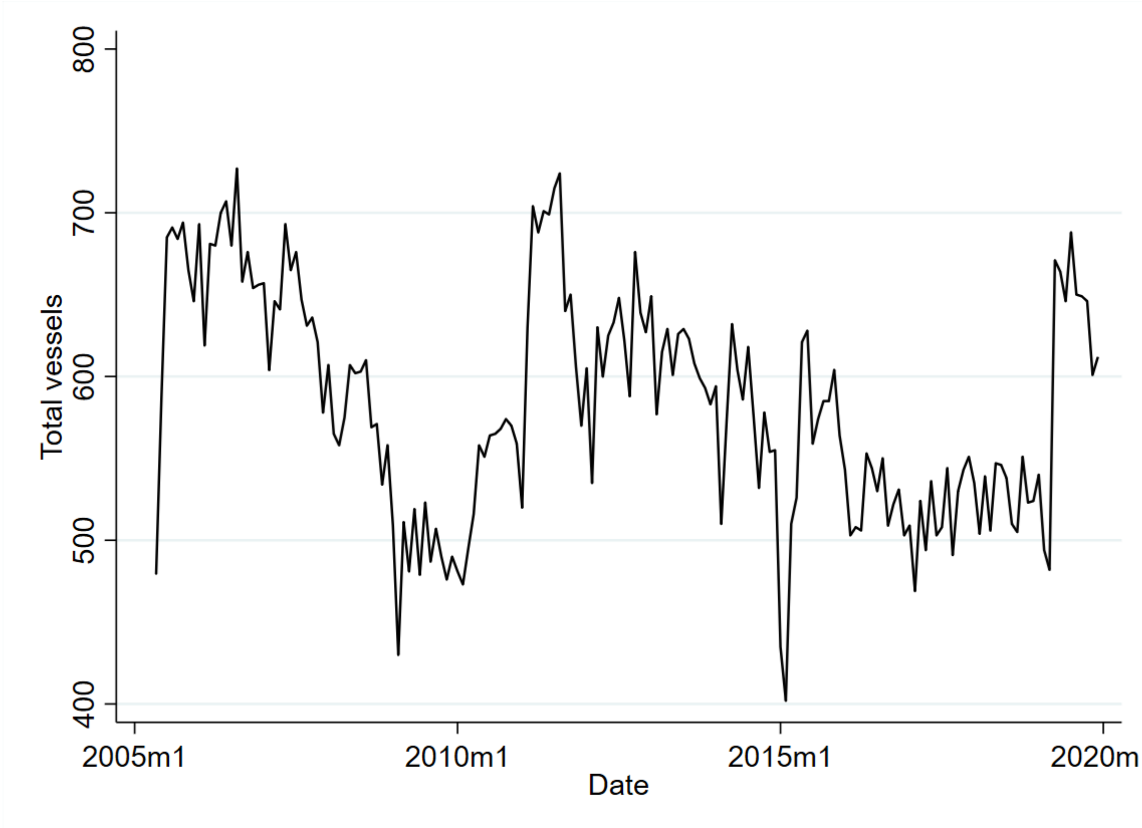

Figure 2.1: Vessels Docking at West Coast

this time. Figure 2.1 shows the total number of vessels docking each month on one of the four ports in the sample. There are many fluctuations, from business cycles, seasonality, and idiosyncratic reasons (like the 2015 work slowdown), but overall there is not a clear trend up or down. A very different graph is Figure 2.2, which shows the average ship length over the same period. Here there is a clear trend upward of 10-15%. I will argue that the economies (or diseconomies) of scale from larger vessels has a major impact on productivity, and failure to consider these changes in the inputs for the ports will bias any productivity measures.

To my knowledge, the dispatch summaries have not been used in any prior research. Gathering and preparing them in a usable form is a contribution of this chapter, so I will briefly describe their construction.

The dispatch summary includes the name of the vessel, how many new gangs to work it, how many gangs are returning, the date and time of the shift, and the date and time of the vessel's arrival. Figure 2.3 shows what each page looks like. Summaries go back to the end of 2004 and up to the present. I limit the sample to spring of 2005 to the end of 2019 because of spotty observations early on and to avoid any impact of the COVID-19

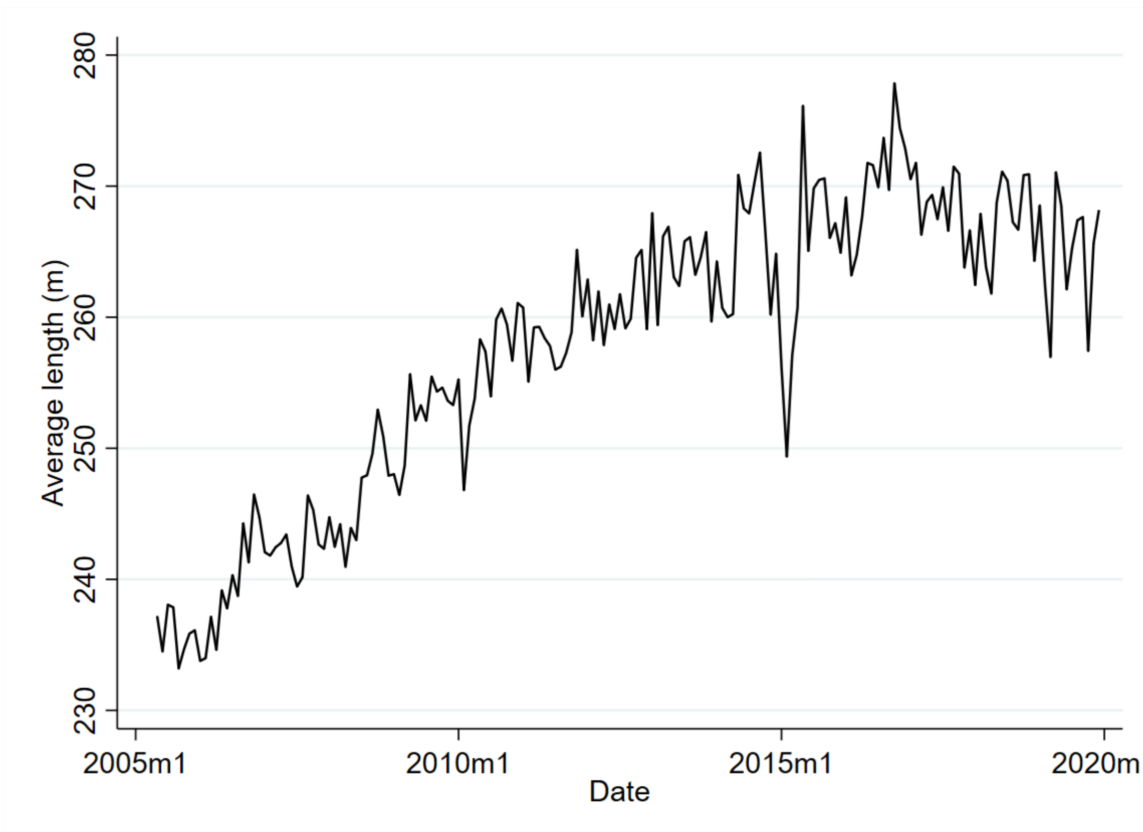

Figure 2.2: Average Vessel Length

	Mean	Standard deviation	Min	Max	Obs
Total gangs	9.70	10.6	1	132	95,620
Total work hours	24.6	20.7	5	513	95,620
Total gang-hours	409	803	5	47,520	95,620
Total hours	44.9	80.4	5	8,532	95,620
Length (m)	252	67.5	6.18	400	63,771
Containers unloaded	386	1,201	1	15,287	38,041

Table 2.2: Summary Statistics

pandemic.

The unit of observation is a stop at a port with at least one hour of work requested, what I call an "unloading." (Some vessels show up for only a few shifts and are idle the whole time. I drop these from the sample.) I identify an unloading by the combination of vessel name, port, and arrival date. When a vessel no longer appears in the summary dispatches, I assume it has left the port. The "total hours in port" are the sum of the shifts the vessel appears in. The "total hours working" are the sum of the shifts where the status is listed as "working." (In Figure 2.3, all the vessels in this shift were being worked on except for the *Azure Sky*, which was idle.) The "total gangs" is the running total `Gangs Worked to Date` in the last shift of the stop. Sometimes this column is blank, in which case I use `Gangs To Complete`. Finally, the "gang-hours" are the product of the total number of gangs and the total number of hours of work.

The summaries list the total number of "gangs" requested. The number of workers on each gang is not listed. If the number or types of workers on a gang were flexible or changing over time, that would bias labor productivity. However, that does not appear to be the case. Contract rules are clear about what kinds of workers compose a gang and how many of them there are, and these rules do not change over time. It is of course possible that as technology changes, these workers are not performing the same tasks. I do not take a stand on whether these particular rules are efficient or not.

2.5 Estimation

Each observation is at the level of unloading. For each estimation, I also run a version with aggregated data at the port and coast level. These versions are identical to the estimations at the level of unloading except for fixed effects that cannot be identified at more aggregated levels.

10/06/2010 2nd Shift LA/LB Dispatch Summary

PORT REGION: Los Angeles / Long Beach
SHIFT: Night
SUMMARY DATE: 10/06/2010
VESSELS:

Company	Vessel	Cargo Type	Status	Gangs Ordered	Gangs Shorted	Gangs Back	Idle Shifts	Gangs Worked to Date	Gangs To Complete	Gangs Unfilled	Arrival Date	Arrival Time
METROPOLITAN STEVEDORE COMPANY	SEA BELL	bulk	working	1	0	0			2		10/03/2010	08:50 PM
PASHA STEVEDORING & TERMINALS L.P.	NASSAU PARADISE	bulk	working	1	0	0		1	6		10/06/2010	07:00 AM
METROPOLITAN STEVEDORE COMPANY	AZURE SKY	bulk	idle-bc	0	0	0		5	5		09/30/2010	01:30 PM
MARINE TERMINALS CORPORATION OF LOS ANGELES	YM GREEN	container	working	0	0	3		19	27		10/05/2010	05:00 AM
YUSEN TERMINALS INC	IKOMA	container	working	3	0	0			3		10/06/2010	12:00 PM
STEVEDORING SERVICES OF AMERICA	MSC TORONTO	container	working	2	0	4		18	37		10/05/2010	04:00 AM
TRAPAC INC	MOL LOIRE	container	working	0	0	2		5	17		10/05/2010	02:00 PM
CALIFORNIA UNITED TERMINALS	HYUNDAI COLOMBO	container	working	1	0	0			27		10/06/2010	04:00 PM
MARINE TERMINALS CORPORATION OF LOS ANGELES	HANJIN DALLAS	container	working	1	0	4		24	35		10/04/2010	03:00 PM
STEVEDORING SERVICES OF AMERICA	COSCO LONG BEACH	container	working	2	0	4		16	25		10/04/2010	04:00 PM
EAGLE MARINE SERVICES LTD.	APL ATLANTA	container	working	4	0	0			20		10/06/2010	03:00 PM
EAGLE MARINE SERVICES LTD.	APL BELGIUM	container	working	2	0	0		4	10		10/06/2010	05:00 AM
MARINE TERMINALS CORPORATION OF LOS ANGELES	XIN DA YANG ZHOU	container	working	4	0	0		5	28		10/06/2010	05:00 AM
YUSEN TERMINALS INC	NYK ARTEMIS	container	working	5	0	0			27		10/06/2010	03:00 PM
EAGLE MARINE SERVICES LTD.	APL ENGLAND	container	working	0	0	4		20	31		10/04/2010	03:00 PM
MARINE TERMINALS CORPORATION OF LOS ANGELES	EVER EXCEL	container	working	1	0	5		25	31		10/04/2010	10:00 PM
STEVEDORING SERVICES OF AMERICA	SUZHOU DRAGON	container	working	5	0	0			10		10/06/2010	02:00 PM
LONG BEACH CONTAINER TERMINAL INC.	OOCL TIANJIN	container	working	0	0	5		21	38		10/04/2010	04:00 PM
TRAPAC INC	MOL PACE	container	working	0	0	3		18	25		10/04/2010	06:30 PM
STEVEDORING SERVICES OF AMERICA	MSC VIENNA	container	working	5	0	0		3	8		10/06/2010	05:00 AM
METROPOLITAN STEVEDORE COMPANY	MATARIKI FOREST	bulk	working	1	0	0			3		10/02/2010	04:00 AM
MARINE TERMINALS CORPORATION OF LOS ANGELES	YM COSMOS	container	working	3	0	0			26		10/06/2010	03:00 PM

Figure 2.3: Screenshot of a dispatch summary

2.5.1 Productivity

My primary measure of productivity is quantity unloaded per gang-hour. Quantity is measured in containers but here a "container" is just a unit of measure equivalent to 17 tons for bulk goods.

To construct this productivity measure for each unloading I need to know the amount of labor used, the size of the vessel, and the amount unloaded. The largest sample includes labor used and vessel size, and a smaller sample includes vessel size and amount unloaded.[4]

I can therefore regress

$$q_{it} = \beta_s s_{it} + PortFE_s + \varepsilon_{s,it} \tag{2.5}$$

directly on the smaller sample and use this to predict q_{it} for all unloadings and to identify η from equation (2.1), $\beta_{s,1}$. Note that the fixed effect here does not vary with time. This is a limitation of this particular sample, which displays plenty of cross-sectional variation but only for a limited number of years. I use the larger sample to regress

$$\ell_{it} = \beta_{prod} s_{it} + Port \times TimeFE_{prod} + \varepsilon_{prod,it} \tag{2.6}$$

which identifies $\frac{\eta-\gamma}{\alpha}$ from (2.3).[5] Because the data are highly heteroskedastic and I am measuring an elasticity, I use Poisson pseudo-maximum likelihood following Silva and Tenreyro (2006) and Sun et al. (2011).

[4] Merging these two into a single dataset is theoretically straightforward. In practice it requires matching not only vessels but vessel-date-port tuples. Small differences in when a vessel is recorded to have docked make the problem messy.

[5] I cannot identify α and γ separately, though this is unimportant for my main question.

2.5.2 Effects of contracts

Let the estimated productivity be $\hat{a}_{it} = \hat{q}_{it} - \hat{\ell}_{it}$, where hats signify predicted values. I am interested in the effects c_y of the new contracts. Because the contract effect is the same for any port to which it applies, I would ideally see both ports affected and unaffected. If productivity trends were the same for all ports, I could measure c_y simply by comparing the differences across ports. However, the contracts are coast-wide, and I therefore need to make additional assumptions. I assume that outside of the effect of the contracts, growth in g_{pt} is constant, $g_{pt} = constant_p + r_p t$. Including a time trend then accounts for these additional effects.

I estimate the regression

$$\hat{a}_{it} = \beta_L L_{it} + \sum_c \mathbb{1}_{\{t \in c\}} \beta_c + \sum_p \mathbb{1}_{\{port=p\}} \Big(\delta_p + r_p \times t\Big) + \varepsilon_{it} \tag{2.7}$$

This is the empirical version of equation (2.4) with the effects of the contracts. The parameters β_c identify the effects of contracts c_y. I also include the logged length of the vessel on the right-hand side for the sake of symmetry with equation (2.4), though recall that $\hat{a}_{it}$ is a deterministic function of L_{it} and therefore the coefficient β_L here does not have any separate meaning; it's just to condition on size.

There are good reasons to be unsatisfied with just including a time trend. First, it is unlikely productivity growth is constant over time, and it is not difficult to think of patterns in the data that could lead to inaccurate estimates. Suppose productivity is roughly constant until a plunge near the end of the period. The time trend would be negative, and the estimated contract effects would be inflated. On the other hand, suppose that all changes in productivity are due to the new contracts, that there is otherwise no other trend. Some of that growth will be attributed to the linear trend, and so I will actually underestimate the effects of the contract.

2.5.3 Effects of automation

I cannot fully address the issues just described because I do not observe a group of ports unaffected by the contracts. However, even though all ports were *allowed* to automate after 2008, only a few terminals actually did. I can, then, compare ports that automated to those that did not while holding fixed any coast-wide time effect. I regress the following

$$\hat{a}_{it} = \beta_{L,auto} L_{it} + \mathbb{1}_{\{it \in auto\}} \beta_{auto} + \sum_{p,t} \mathbb{1}_{\{port=p, time=t\}} \delta_{pt} + \varepsilon_{it} \tag{2.8}$$

Here, rather than a linear time trend, there are separate port-month fixed effects allowed to vary flexibly. The indicator $\mathbb{1}_{\{it \in auto\}}$ is positive only for ports with a terminal that automated. Only one port had such a terminal, and it only automated in 2016, so I can estimate the differential impact of automation with more confidence.

I should emphasize here that (2.8) is not a better version of (2.7). The equations are measuring different effects. It is possible that even among ports that did not build fully automated terminals, the new terms led to changes in productivity. Perhaps terminal operators were more likely to make technical investments short of fully automation, being less uncertain about potential union responses. Workers may also have behaved differently, knowing future automation was a real possibility. As often happens, there is one effect that I am most interested in knowing and which is hard to reliably estimate. There is another which is easier to accurately estimate and is of less interest. In this chapter I report both.

2.6 Results

2.6.1 Productivity

Tables 2.3 and 2.4 show the relationship between quantity unloaded and length and labor employed and size, respectively. The leftmost column of each table includes the fixed effects described in the Estimation section. In the next two columns of each, I aggregate up to the port level and then the entire coast level and perform similar regressions. I say "similar" because of course more aggregated versions are not necessarily identified with port-time or even port-level fixed effects. I also show versions of Unloading and Port without any fixed effects to make the regressions more comparable; these are the last two columns of each.

The quantity unloaded roughly increases with the square of the length of the ship. This makes sense: if quantity moves one-to-one with volume and ships differed only by length, a ship twice as long would hold twice as much and we would expect a linear relationship between length and quantity. If a twice as long ships was also twice as wide and had a draft twice as high, we would expect a cubic relationship. The reality is likely somewhere in between, which is what we estimate.

The amount of labor necessary increases at a faster rate, 2.76. That β_{prod} is greater than β_s is important: it means that the additional output a larger vessel brings requires more than a proportional amount of labor to unload; there are diseconomies of scale.

In Figure 2.4 I plot the containers per gang-hour using the coast-wide aggregate labor productivity and the productivity of the average unloading. This shows levels, rather than the changes I am estimating in the regressions, but already it is clear that 1) aggregate and average persistently differ and 2) the amount by which they differ varies over time.

Aggregation has large effects on the estimated parameters.[6] In Table 2.3, an estimation

[6] The variables used here are the logged total number of containers, logged total gang-hours, and logged sum of the lengths of all unloadings. If I use logged mean lengths I get similar results.

even at the level of port-month produces an elasticity of about zero. Table 2.4 is even more dramatic, there the sign of β_{prod} actually flips at the port level.

	Unloading	Port
β_s	2.08***	0.0741
	(0.119)	(0.0540)
Fixed Effects?	Y	Y
N	31,182	72

Table 2.3: Containers unloaded as function of ship length

*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Fixed effects are port fixed effects for columns one and two.

	Unloading	Port
β_{prod}	2.76***	-2.90*
	(0.0639)	(1.45)
Fixed Effects?	Y	Y
N	62,055	703

Table 2.4: Labor employed as function of ship length

*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Fixed effects are port × month for first column and separate port and time for second.

2.6.2 Effects of contracts

Column 1 of Table 2.5 shows the main results of the contracts. Removing the trend, the 2008-2015 contract increased labor productivity by about 25%. The productivity falls during the regime of the third contract to only 15% above the pre-2008 level. Though I include standard errors in this table as in previous ones, these have not been adjusted for the fact that the left-hand side of the regressions is itself an estimator.

The fall in the productivity during the 2015-2019 contract may appear puzzling. The time trend (not shown) is negative, which means $\beta_{2015-2019}$ is, if anything, likely to be biased up. There was nothing in the 2015 contract that retracted the automation clause in the 2008 one. In fact, actual automation did not begin in earnest until 2016, when the 2015 contract was in effect. What could explain the drop?

One possibility is that, as noted earlier, though the 2015 contract did not include many novel clauses, the negotiation process surrounding it was contentious. The fact that it was a a

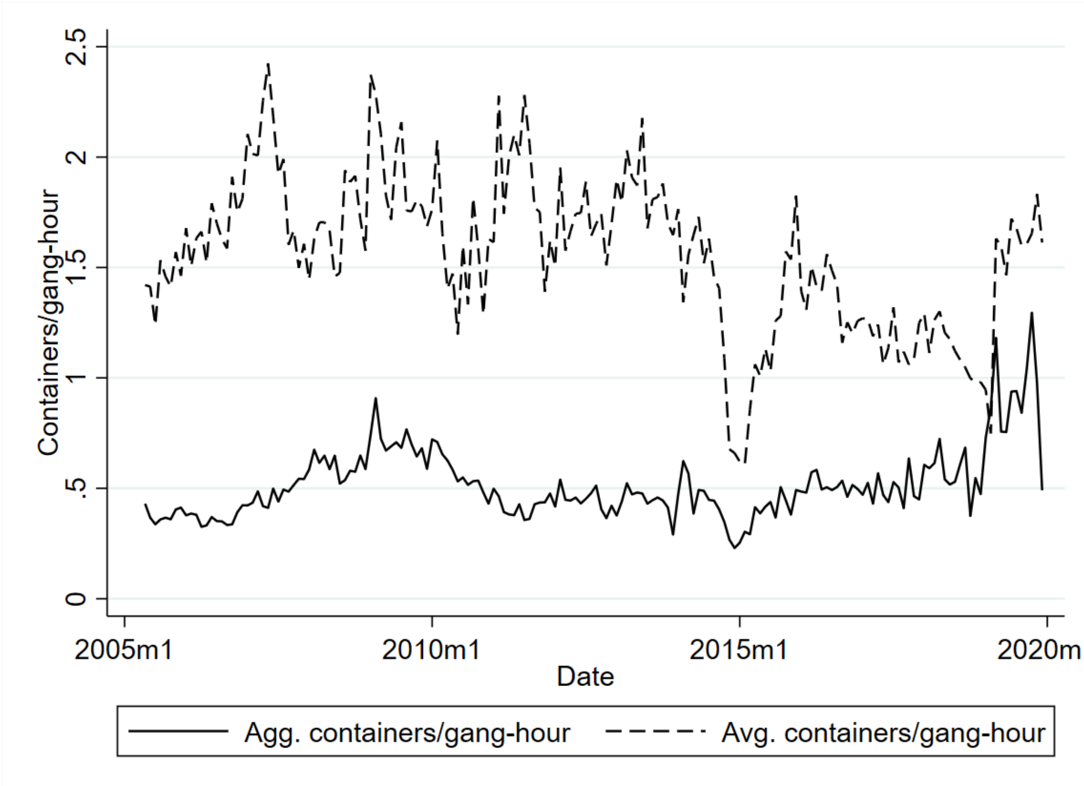

Figure 2.4: Coast-wide versus average unloading productivity

contract signed in 2015 rather than when the old one expired in 2014 highlights how drawn-out the negotiating process was compared to 2008. The actual work slowdown in early 2015 does not explain the difference; that occurred in the previous contract regime. However, if after the 2008 contract the ILWU appeared surprisingly cowed, the 2015 slowdown made clear that was not the case. Employers who previously may have expected less push back from automation or further changes to work rules could have updated their beliefs, as would have the workers themselves. In fact there has been almost a cyclical nature to negotiations: after a lockout in 2002, there was little drama in 2008, followed by the slowdown in 2014-2015 Mongelluzzo (2015).

Table 2.5 also shows results for port and coast aggregations. As with the earlier regressions, the results are very different from the unloading. At the coast level, there is basically no effect. At the level of port, the effect is much larger and the relative order is flipped, $\beta_{2015-2019}$ is actually greater than $\beta_{2008-2015}$.

At this point I should emphasize again that the regressions in the second and third columns of Table 2.5 are not wrong or meaningless. The port-level labor productivity is a well-defined object. "Aggregate length" is maybe less well-defined (or at least objections could be raised to the way in which I am aggregating it), but a notion of total capacity is certainly coherent.

However, as I am studying the effect on the technology, on the process of unloading itself, these are not the numbers I am interested in.

At level of...	Granular (Unloading)	Aggregate (Port)
$\beta_{2008-2015}$	0.259***	0.874***
	(0.00629)	(0.107)
$\beta_{2015-2019}$	0.146***	1.30***
	(0.0126)	(0.217)
N	62,056	466

Table 2.5: Contract effects on logged containers per gang-hour

*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Not shown: Time trend and port fixed effects.

2.6.3 Effects of automation

Table 2.6 shows the results of the difference-in-difference with the Long Beach Container Terminal becoming fully automated in April 2016. The containers per gang-hour productivity increased by around 15% compared to the other ports at the same time. This is a sizeable increase, but it's less than a quarter of the estimated effect using data aggregated at the port level.

	Granular (Unloading)	Aggregate (Port)
β_{auto}	0.146*	0.629*
	(0.0676)	(0.264)
Port FE?	Y	Y
Time FE?	Y	Y
N	62,056	466

Table 2.6: Long Beach automation difference-in-difference

*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Standard errors are clustered at the port level.

Figure 2.5 shows the regression visually. The points are the differences from the mean productivity conditional on port, time, and ship size. Productivity in Long Beach actually fell, it just fell less than at all other ports.

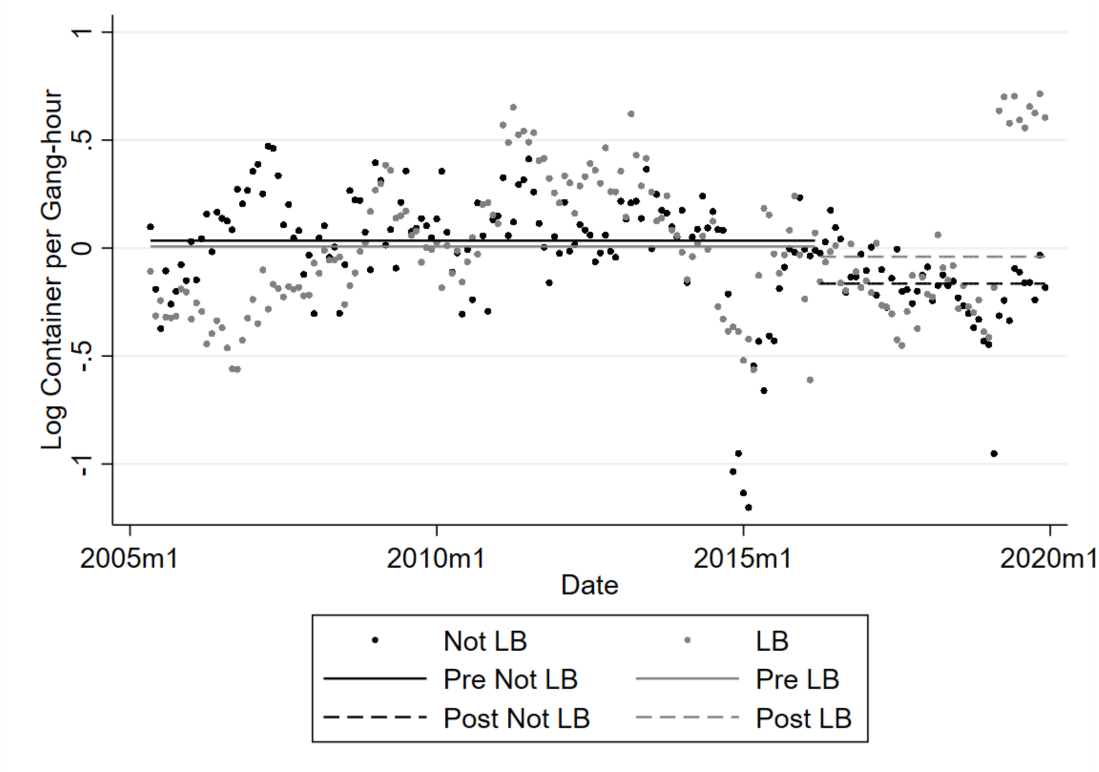

Figure 2.5: Gang-hours before and after LBCT automation

2.7 Robustness

I extrapolate quantity unloaded–and through that, productivity–based on bill of ladings data that is incomplete in two ways. First, the data do not cover the full time period but only a handful of months in 2007, 2008, and 2012-2015.[7] These dates do not include a post-automation time period, but I can run a version of the estimation of contract effects using only these dates. Second, the bill of ladings only includes containerships. Ships are specialized for different kinds of cargo and the relationship between quantity and size may not be constant across different types. I therefore re-estimate the effects of contracts using only containerships.

Table 2.7 shows the effects of the 2008 contract in an estimation limited to only the months were I observe quantity unloaded directly. I cannot estimate $\beta_{2015-2019}$ because these dates do not go beyond March 2015. The aggregate estimate is still higher than the granular, though they are closer and not statistically significantly different. At least part of this is because the sample sizes are a tenth of the full sample. There is also only one month

[7] The specific months covered are December 2007; November and December 2008; November and December 2012; January, February, March, November, and December 2013, January, February, March, November, and December 2014; and January, February, and March 2015.

At level of...	Granular (Unloading)	Aggregate (Port)
$\beta_{2008-2015}$	0.566***	0.616*
	(0.0154)	(0.306)
N	6,227	49

Table 2.7: Contract effects on logged containers per gang-hour: restricted dates
*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Not shown: Time trend and port fixed effects.

outside of the 2008 contract regime in this sample, so the results are highly dependent on that specific time. Nevertheless, it is encouraging that the orders of magnitude and signs are not different across specifications.

At level of...	Granular (Unloading)	Aggregate (Port)
$\beta_{2008-2015}$	0.348***	0.523***
	(0.00520)	(0.0854)
$\beta_{2015-2019}$	0.281***	0.745**
	(0.0113)	(0.195)
N	52,291	454

Table 2.8: Contract effects on logged containers per gang-hour: containerships
*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Not shown: Time trend and port fixed effects.

In Table 2.8, I show the results of estimating the contract and automation effects, respectively, with only containerships in the sample. Notice that the total number of observations is not that much lower than the full sample: the vast majority of unloadings at these ports are from containerships. The patterns are largely the same as with the full sample, though as in the sample restricted to certain dates, the estimates values for the aggregate and granular versions are closer. As the ships in the sample become more homogeneous, it makes sense that aggregation loses less information.

The granular results for both regressions may be biased downwards if there is measurement error in the independent variables. In Appendix B I discuss where this error may occur and how it compares to the aggregation bias.

2.8 Conclusion

Far beyond just determining pay or labor supplied, unions can determine the technology of production itself. In this chapter, I have shown how an unprecedented change in the ILWU contract led to productivity growth of 25%. Part of this was surely the automation itself, but there was also growth before any of the terminals were actually automated. In fact, the period of time when the terminals were automated was actually one of less growth, though the terminal itself had a positive effect. Beyond the formal rules of what is allowed and disallowed, the general attitude the union and employers take towards one another likely affects day-to-day work and productivity. If the ILWU decides it has been taken advantage of in recent negotiations, productivity may suffer even as the set of technically allowed investments expands.

Bibliography

American Association of Port Authorities. Results of AAPA's Port Planned Infrastructure Investment Survey, 2016.

American Association of Port Authorities. Glossary of maritime terms, 2019. URL https://www.aapa-ports.org/advocating/content.aspx?ItemNumber=21500. Website accessed January 29, 2020.

Army Corps of Engineers. Discount Rates in the Economic Evaluation of U.S. Army Corps of Engineers Projects. Technical report, Army Corps of Engineers, 2007.

Army Corps of Engineers. U.S. Port and Inland Waterways Modernization: Preparing for Post-Panamax Vessels. Technical report, US Army Corps of Engineers: Institute for Water Resources, June 2012.

Samuel Bailey and Thomas J. Holmes. Scale economies and alliances in ocean shipping. Working paper, in progress.

Timothy J. Bartik. Should Place-Based Jobs Policies Be Used to Help Distressed Communities? Technical report, W.E. Upjohn Institute, August 2019.

Susanto Basu and John G Fernald. Aggregate productivity and aggregate technology. *European Economic Review*, page 29, 2002.

Steven T Berry and Joel Waldfogel. Free Entry and Social Inefficiency in Radio Broadcasting. *RAND Journal of Economics*, 30(3):397–420, 1999.

William Booth. Expanded Panama Canal sparks race to be ready for bigger cargo ships. *Washington Post*, 2013. ISSN 0190-8286.

Benjamin Bridgman. Why Containerization Did Not Reduce Ocean Trade Shipping Costs. Working paper, October 2018.

Matthew Chambers and Mindy Liu. Maritime Trade and Transportation by the Numbers. Technical report, Bureau of Transportation Statistics, 2012.

Kevin Cullinane and Mahim Khanna. Economies of Scale in Large Container Ships. *Journal of Transport Economics and Policy*, 33(2):185–207, 1999. ISSN 0022-5258.

Peter W. De Langen and Athanasios A. Pallis. Analysis of the Benefits of Intra-Port Competition. *International Journal of Transport Economics*, 33(1):18, February 2006.

WCN Editorial. ILWU fights grant funding for automated equipment. https://www.worldcargonews.com/news/news/ilwu-fights-grant-funding-for-automated-equipment-65625, 2021.

Jasper Faber, Maarten 't Hoen, Robert Vergeer, and John Calleya. Historical Trends in Ship Design Efficiency. Technical Report 16.7H27.23, CE Delft, March 2016.

Finance and Administration Bureau of Port of Los Angeles. 2018 comprehensive annual financial report, 2019. Annual report.

Austan Goolsbee and Amil Petrin. The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV. *Econometrica*, 72(2):351–381, March 2004. ISSN 0012-9682, 1468-0262. doi: 10.1111/j.1468-0262.2004.00494.x.

Zvi Griliches. The Sources of Measured Productivity Growth: United States Agriculture, 1940-60. *Journal of Political Economy*, 71(4):331–346, August 1963. ISSN 0022-3808, 1537-534X. doi: 10.1086/258782.

Thomas J. Holmes and James A. Schmitz. Competition at Work: Railroads vs. Monopoly in the U.S. Shipping Industry | Federal Reserve Bank of Minneapolis. *Quarterly Review*, 25(2):3–29, 2001. doi: 10.21034/qr.2521.

Thomas J Holmes and Ethan Singer. Indivisibilities in Distribution. Working Paper 24525, National Bureau of Economic Research, April 2018.

Charles R. Hulten and Robert M. Schwab. A fiscal federalism approach to infrastructure policy. *Regional Science and Urban Economics*, 27(2):139–159, April 1997. ISSN 01660462. doi: 10.1016/S0166-0462(96)02150-3.

Masahiro Ishii, Paul Tae-Woo Lee, Koichiro Tezuka, and Young-Tae Chang. A game theoretical analysis of port competition. *Transportation Research Part E: Logistics and Transportation Review*, 49(1):92–106, January 2013. ISSN 13665545. doi: 10.1016/j.tre.2012.07.007.

Jan Owen Jansson and Dan Shneerson. Economies of Scale of General Cargo Ships. *The Review of Economics and Statistics*, 60(2):287, April 1978. ISSN 00346535. doi: 10.2307/1924982.

Joint Legislative Audit and Review Commission. Review of the Virginia port authority's competitiveness, funding, and governance. Technical report, Joint Legislative Audit and Review Commission, 2013.

Chung-Yee Lee and Qiang Meng, editors. *Handbook of Ocean Container Transport Logistics: Making Global Supply Chains Effective*, volume 220 of *International Series in Operations Research & Management Science*. Springer International Publishing, Cham, 2015. ISBN 978-3-319-11890-1 978-3-319-11891-8. doi: 10.1007/978-3-319-11891-8.

Arthur Lewbel. Aggregation with Log-Linear Models. *The Review of Economic Studies*, 59 (3):635, July 1992. ISSN 00346527. doi: 10.2307/2297869.

N. Gregory Mankiw and Michael D. Whinston. Free Entry and Social Inefficiency. *The RAND Journal of Economics*, 17(1):48–58, 1986. ISSN 07416261. doi: 10.2307/2555627.

Olaf Merk, Bénédicte Busquet, and Raimonds Aronietis. The Impact of Mega-Ships. Technical report, International Transport Forum, 2015.

Bill Mongelluzzo. ILWU-PMA contract no game changer for West Coast productivity | JOC.com. https://www.joc.com/port-news/longshoreman-labor/international-longshore-and-warehouse-union/ilwu-pma-contract-no-gamechanger-west-coast-productivity_20150523.html, 2015.

Congressional Budget Office. Public Spending on Transportation and Water Infrastructure, 1956 to 2017, 2018.

Ariel Pakes. On Group Effects and Errors in Variables in Aggregation. *The Review of Economics and Statistics*, 65(1):168, February 1983. ISSN 00346535. doi: 10.2307/1924426.

Valerie Ramey. The Macroeconomic Consequences of Infrastructure Investment. Technical Report w27625, National Bureau of Economic Research, Cambridge, MA, July 2020.

Dan Ronan. Bayonne Bridge Expansion Pays Big Dividends for Port of N.Y.-N.J. https://www.ttnews.com/articles/bayonne-bridge-expansion-pays-big-dividends-port-ny-nj, March 2020.

James A. Schmitz. What Determines Productivity? Lessons from the Dramatic Recovery of the U.S. and Canadian Iron Ore Industries Following Their Early 1980s Crisis. *Journal of Political Economy*, 113(3):582–625, 2005.

Rexford B. Sherman. Seaport Governance in the United States and Canada. Technical report, American Association of Port Authorities, 2008.

J. M. C. Santos Silva and Silvana Tenreyro. The Log of Gravity. *Review of Economics and Statistics*, 88(4):641–658, November 2006. ISSN 0034-6535, 1530-9142. doi: 10.1162/rest.88.4.641.

Cailin Slattery. Bidding for Firms. Working paper, 2019.

Kai Sun, Daniel J. Henderson, and Subal C. Kumbhakar. Biases in approximating log production. *Journal of Applied Econometrics*, 26(4):708–714, June 2011. ISSN 08837252. doi: 10.1002/jae.1229.

Charles M. Tiebout. A Pure Theory of Local Expenditures. *Journal of Political Economy*, 64(5):416–424, October 1956. ISSN 0022-3808, 1537-534X. doi: 10.1086/257839.

Peter Tirschwell. Longshore labor: ILWU's anti-automation plea to California hints at contract flashpoint. https://www.joc.com/port-news/us-ports/ilwu%E2%80%99s-anti-automation-plea-california-hints-contract-flashpoint_20200724.html, 2020.

Matthew Turner, Gilles Duranton, and Geetika Nagpal. Transportation Infrastructure in the US. Technical Report w27254, National Bureau of Economic Research, Cambridge, MA, May 2020.

Appendix A

Appendix to Chapter 1

A.1 Charter rates

Drewry Insights publishes average daily charter rates for container ships of various sizes. These rates do not include fuel but include labor, insurance, and other operational costs. Most carriers own most of their own fleets, but most also charter at least some ships and so it is reasonable to assume the charter rates cannot vary much from the cost of operating an owned vessel (they certainly cannot be much lower).

I observe 43 observations from 2004 to 2018. The sizes listed are 500, 700, 3,500, 4,250, 5,000, and 8,500 TEU. I estimate

$$\text{log\_charter\_rate}_i = \text{constant} + (1 - \gamma_{charter}) \, \text{log\_TEU}_i + \sum_{t=2004}^{2018} \tau_t \mathbb{1}_{year_i = t} + \varepsilon_i$$

(Jansson and Shneerson (1978)) and (Cullinane and Khanna (1999)) use an annualized capital cost by regressing the price of newly built ships on size. I use charter rates rather than newbuilds as I do not have the data on costs like crew that these other papers use for my main results. However, I perform a similar regression to compare my number to past ones, and get very similar results. (See Table 1.7.)

A.2 Fuel costs

In their paper on hull design efficiency, (Faber et al. (2016)) use the following relationship between power, speed, and physical dimensions:

$$\log P_{ME} = c \times \log (V \times \text{frictional resistance})$$

where c is a constant, V is speed, and "frictional resistance" is a complicated expression that depends on speed, length, draft, and beam. I know that the power required to move the Maersk ship *Emma* is 80,800 kilowatts and I use this to get the correct constant c. For the speed I use the design speed. Finally, the relationship between fuel consumption and power is

$$FC = SFOC \times P_{ME}$$

where FC is fuel consumption and $SFOC$ is the specific fuel consumption of the engine. (Faber et al. (2016)) use a value of 190, though (Cullinane and Khanna (1999)) use the much lower 125. I use 170. This value is not unimportant as it affects the total fuel usage linearly, but the final results do not change much if I vary it from 100 to 200. (At 100 the final γ is 0.34 rather than 0.28; at 200 it is 0.26.)

A.3 Total cost

To get total costs, I convert the fuel usage to daily fuel usage assuming 24 hours of movement. I take the cost of fuel to be $493, which was the average Brent oil price from 2002 to 2019. (Bunker fuel and crude oil prices are close and move together.) I add the predicted fuel costs to predicted charter costs, take the log, and regress on logged TEU,

$$\log\left(\text{predicted charter}_i + \text{predicted fuel}_i\right) = constant + (1-\gamma)\log TEU + \varepsilon_i$$

Appendix B

Appendix to Chapter 2

B.1 Aggregation bias or measurement error?

If observations are measured with error, the granular regressions will be biased towards zero and may be no more reliable than the aggregated regressions. In fact, for classical measurement error, aggregating by groups may lead to more consistent results ((Pakes, 1983)).[1]

There are two potentially important sources of measurement error in the data I use. First, I observe hours at the level of a shift. If a ship is present and being worked on at the beginning of a shift, I do not observe if it stays the whole shift, leaves halfway through, or leaves one hour after the shift starts. A second source of error are worker types. There is a rigid seniority list in who is called up for work. In periods of less work, the only people called will be senior union members, who have longer tenure and more experience. When there is more work the employers may call up people from the casuals roll, who have much less experience. I treat these workers as identical, potentially erasing important differences in productivity.

That being said, there are good reasons to think there is also aggregation bias. Arthur Lewbel has shown that for log-linear models, a sufficient and necessary condition for aggregated regressions to return the same parameters as granular regressions is that the distribution of the right-hand side variables are "mean scaled" ((Lewbel, 1992)). In other words, if ship lengths are x_{ipt} and the mean ship length at each port-year-month level is X_{pt}, then the distribution of a variable $z_{ipt} \equiv \frac{x_{ipt}}{X_{pt}}$ should be independent of X_{pt}.[2]

In Figure B.1, I show the relationship between the mean vessel size and another moment of its distribution, standard deviation. Clearly, the distribution of size is not independent

[1] I thank Amil Petrin for pointing this out.

[2] An intuitive example from Lewbel is personal income. Personal income is mean scaled if the Gini coefficient does not change with mean income.

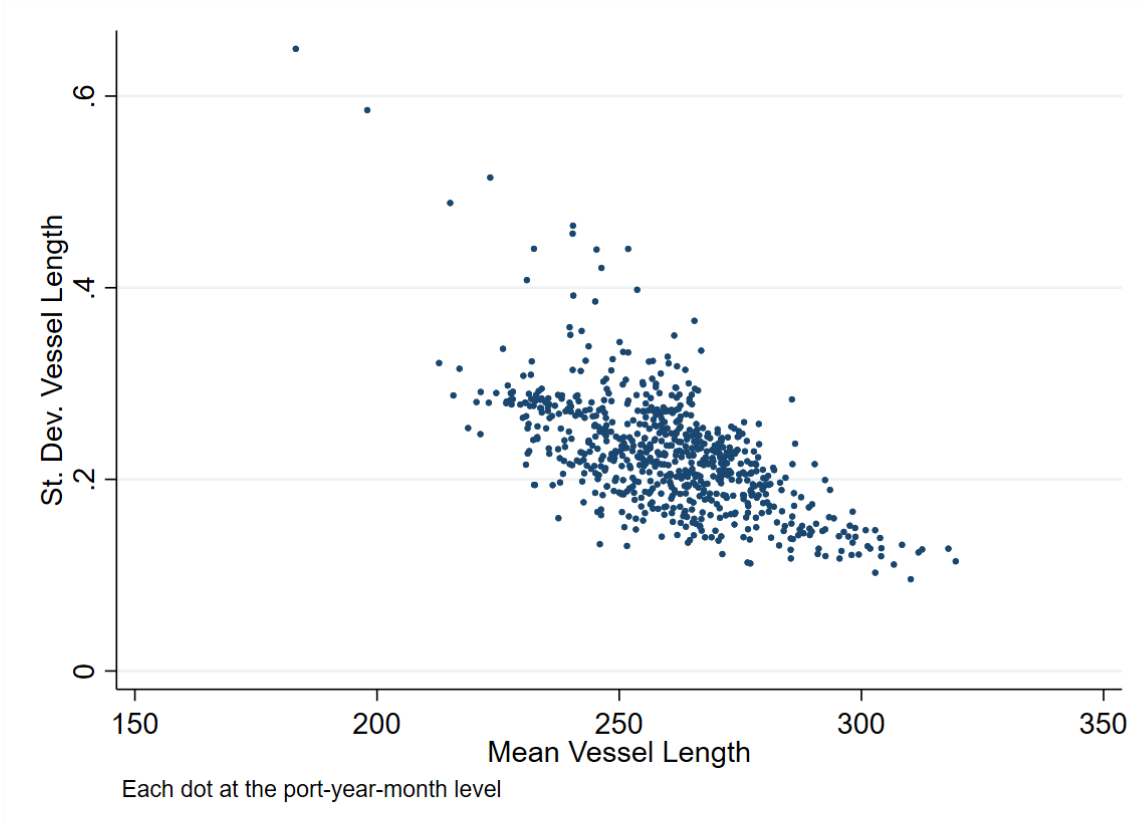

Figure B.1: Mean vessel length versus standard deviation

of the mean size. Port-year-months with larger vessels on average also have less dispersion. Unfortunately it is difficult to say the direction of the bias without knowing the value of the true parameter, but in the case of log normally distributed variables the magnitude increases in proportion to the covariance between $logX_{pt}$ and standard deviation. The granular results are therefore at the very least a lower bound, but it is not clear whether the aggregate results are above or below the true values.