Investment in Infrastructure and Trade: The Case of Ports

Giulia Brancaccio[*], Myrto Kalouptsidi[†], Theodore Papageorgiou[‡§]

New York University, Harvard University, Boston College

May 2024

Transportation infrastructure is vital for the smooth functioning of international trade. Ports are a crucial gateway to this system: with more than 80% of trade carried by ships, they shape trade costs, and it is critical that they operate efficiently. Yet ports are susceptible to disruptions, causing costly delays. With enormous budgets spent on infrastructure to alleviate these costs, a key policy question emerges: in a world with high volatility, what are the returns to investing in infrastructure? To address this question, we introduce an empirical framework that combines insights from queueing theory to capture port technology, with tools from demand estimation. We use our framework, together with a collection of novel datasets, to quantify the costs of disruptions and evaluate transportation infrastructure investment. Our analysis unveils three policy-relevant messages: (i) investing in port infrastructure can lead to substantial trade and welfare gains, but only if targeted properly– in fact, net of costs, investment has positive returns at a minority of US ports; (ii) there are sizable spillovers across ports, as investing in one port can decongest a wider set of ports, suggesting that decision-making should not be decentralized to local authorities; (iii) macroeconomic volatility can drastically change returns to investment and their geography.

Keywords: transportation, infrastructure, ports, congestion, macroeconomic volatility, disruptions, spillovers, welfare

[*]New York University, Stern School of Business, Kaufman Management Center 44 West 4th Street New York NY 10012, USA, giulia.brancaccio@nyu.edu

[†]Harvard University, Department of Economics, Littauer Center 124, Cambridge MA 02138, USA, myrto@g.harvard.edu

[‡]Boston College, Department of Economics, Maloney Hall 343, Chestnut Hill, MA, 02467, USA, theodore.papageorgiou@bc.edu

[§]We are thankful to George Alessandria, Jim Anderson, Costas Arkolakis, Nick Buchholz, Arnaud Costinot, Dave Donaldson, Ed Glaeser, Paul Grieco, Gabriel Kreindler, Isabela Manelici, Marc Melitz, Kostas Metaxoglou, Amedeo Odoni, Ariel Pakes, Steve Redding, Marta Santamaria, and Edouard Schaal for their many helpful comments. We are extremely grateful to Constanza Abuin, Pedro Degiovanni, Toren Fronsdal, Peleg Samuels, Zachary Saunders, and Yixin Zhou for excellent research assistance. We also gratefully acknowledge financial assistance from the NBER Transportation Economics in the 21st Century Grant, as well as NSF Award SES-231533-231534-231535.

1 Introduction

Transportation infrastructure is crucial for the smooth functioning of international trade. As such, enormous funds are spent every year to improve the efficiency and resilience of the infrastructure network. Ports, which are the focus of this paper, are a vital link of this network. With more than 80% of trade (amounting to 11 billion tons and about $20 trillion) carried by ships, they are the literal gateway to international trade. Ports are therefore crucial determinants of trade costs, as became painfully apparent during the disruptions induced by the Covid-19 pandemic, which upended supply chains, and generated enormous costs to importers and exporters worldwide, who faced massive delays in obtaining their goods.

Investment in transportation infrastructure has the potential to alleviate these costs. A key policy question therefore is: in a world with high volatility, what are the returns to investing in infrastructure? To answer this question, we first assemble a collection of novel micro-data on various aspects of port operations. We then introduce an empirical framework that combines insights from queueing theory to capture port technology, with tools from demand estimation. We use this framework to quantify the cost of disruptions and to evaluate transportation infrastructure investment. While our analysis focuses on maritime ports, in principle this framework can be applied to other setups with congestion, such as airports.

Our analysis unveils three policy-relevant messages: (i) investing in port infrastructure can lead to substantial trade and welfare gains, but only if targeted properly– in fact, net of costs, investment has positive returns at a minority of US ports; (ii) there are sizable spillovers across ports, as investing in one port can decongest a wider set of ports, suggesting that decision-making should not be decentralized to local authorities; (iii) macroeconomic volatility can drastically change the returns to investment and their geography.

More specifically, we leverage detailed micro-data on various aspects of port operations, obtained from a number of different sources. Using the universe of port calls, we collect data on port "outputs", which in our context includes port service times and queues, as well as on port demand. In addition, using satellite images, we collect information on port "inputs", namely existing port infrastructure such as berth, storage space and cranes, as well as port employment from union membership. We also obtain access to micro-data on the prices ports charge for their services. Finally, we collect information on the cost of port expansion, such as land values and dredging costs. Our study focuses on US ports and on dry bulkers, which carry about half of the volume of maritime trade. Dry bulk ships are the main mode

of transportation for commodities, such as grain, ore, and coal.

Our key measure of port efficiency is the total time a ship spends at port, which includes both its service time, i.e. the time it takes to discharge or load cargo, but also the time it spends outside the port waiting to be serviced. Ships spend on average 117 hours per trip at US ports. This amounts to roughly half of the average travel time, suggesting that iceberg transport costs, that are a function of distance, miss a key component of trade costs. Time at port can be very costly, both because of the user cost of ships, but also because of any indirect costs of delays along the entire supply chain. As we describe below, our empirical model allows us to measure this cost of port disruptions.

In order to evaluate the role of transportation infrastructure we need the following ingredients: first, we need to accurately capture infrastructure technology, which we do through a port production function. This helps us understand both the dynamics of congestion, as well as how investment in infrastructure affects congestion. Second, we need a demand system for transportation services. This is necessary both because demand endogenously responds to infrastructure improvements, but also in order to quantify their welfare gains. We briefly discuss each ingredient in turn.

Using insights from queueing theory, we propose a micro-model of port technology that takes port inputs, such as infrastructure, labor, and cranes, and produces output, namely time at port. We show that congestion can be highly sensitive to even small demand shocks, making ports susceptible to disruptions. To see this, note that at low levels of demand, time at port is almost flat, reflecting mostly service time; as demand keeps rising, time at port starts increasing and after a certain point port infrastructure can barely keep up. Now even small increases in demand can make the queue explode. In other words, our setup endogenously generates convex costs of congestion. Moreover, increases in volatility exacerbate the fragility of the port system, by making extreme demand realizations more frequent. Investment in infrastructure, by smoothing these fluctuations, can generate substantial welfare gains.

We then estimate a demand system for port services: potential exporter-importer pairs decide which port to use and value lower time at port, shorter distance to their origin and destination, lower port prices, as well as other characteristics of the port. Importantly, we allow for aggregate macroeconomic conditions to shift demand over time. One complication in our analysis is that the ship movement data that we use does not contain information on the final US destination/origin of a good. For example, we may observe a ship load at the port of Houston (TX) and head to Europe, but we do not observe the origin of the shipment within the US. To address this, we augment our analysis with aggregate trade data on freight movements within the US. We estimate the demand system following the approach of Grieco

et al. (2022).

Our estimates suggest that agents are quite sensitive to changes in time at port. Moreover we find that the average ship is willing to pay $45,000 to reduce the time at port by one day. This amount is substantially higher than the direct user cost of the ship, thus capturing the cost of port disruptions as they propagate through the supply chain. Although the geographic and other characteristics of the port do affect demand, we find that 83% of shippers do not choose the most geographically convenient port. This suggests that time at port acts as a barrier to trade, forcing shippers to sacrifice geographic convenience in order to avoid port congestion. Finally, substitution patterns vary greatly by port: for instance, ports in the Great Lakes substitute only locally, while ports in the Gulf and East Coast boast a wide range of substitutes as they serve a much wider set of origins and destinations.

Having all necessary ingredients, we combine our estimates of technology with our estimates for port demand to estimate the welfare gains of port infrastructure investment. We do so by considering the following question: if we could increase the capacity of the US port system by one slot, i.e. a single port can now handle one additional ship, where would we allocate this additional slot? To put this in context, this increase corresponds to approximately a 1% increase of the total US capacity.

We find that this increase in effective port capacity increases US trade and welfare by as much as 1.3% and 0.8% respectively, but that the effect is quite heterogeneous across ports. The highest welfare gains are obtained for ports in the East Coast and the Gulf, such as Hampton Roads (VA), Houston (TX), and Baltimore (MD). At the other extreme, investing in the Great Lakes barely makes a dent in total trade and aggregate welfare. This geographical dispersion in the returns from investment showcases that targeting can increase the gains from investment manyfold.

What makes a port generate high (vs. low) welfare gains? When capacity increases at a port, demand there spikes from an inflow of both new shippers, as well as those that substitute away from other ports.[1] The more geographically attractive, and the more congested the port, the higher the increase in its demand. For instance, Hampton Roads (VA), the port with the highest returns, is able to serve many and diverse origins and destinations, which lowers the transportation costs that shippers need to shoulder.

But that is not all. In this inter-connected system, changing capacity in one port does not happen in a vacuum. Expanding capacity at a port reduces demand at substitute ports, which now become decongested. In other words, investment at one port has positive spillovers to other ports, which in our setup explain up to 39% of the final aggregate increase in trade for the ports with the highest gains.

[1]The endogenous response of demand dampens the original fall in congestion. This is reminiscent of the "Iron Law of Congestion" (Downs, 1962); the increase in trade however, generates positive welfare gains.

Welfare gains thus do not just depend on the attributes of the treated port, but also on which other ports it decongests. Hampton Roads (VA) decongests a large set of ports that are themselves geographically attractive and congested, amplifying its welfare gains. The presence of these spillovers raises important policy questions regarding the coordination of investment activities and suggests that unless decisions are made centrally, they are unlikely to internalize the spillovers across ports.

Finally, when and where is port investment justified? Can the increasing macroeconomic volatility warrant more investment? To answer these questions, we bring in cost estimates from the US Army Corps of Engineers, which are used by authorities to make investment decisions, in order to compute net returns of investment; i.e. the difference between welfare gains and infrastructure costs. We find that net returns are positive for 15 out of 51 US ports, and are highest for the ports in the Gulf, such as Corpus Christi (TX) and Convent (LA), as well as for Hampton Roads (VA).

Next, in line with concerns about disruptive shocks becoming more common (see Bloom, 2009, as well as the literature on the recent supply chain disruptions), we compute how frequent macroeconomic shocks need to be to warrant investment. An increase in volatility makes both very large and very low demand realizations more likely. While low demand realizations matter little, high demand shocks can lead to sharp increases in congestion. As a result, higher volatility makes infrastructure investment more valuable: on average, welfare gains double when volatility doubles, and net returns turn positive for 10 additional ports. At one extreme, welfare gains quadruple for some ports. However, at the other extreme, for a sizable share of ports (25%) net returns do not turn positive even when volatility more than triples; for these ports, investment is unlikely to be warranted even if the policymaker only cares about resilience.

In addition, higher macroeconomic volatility changes the geography of the returns. Interestingly, infrastructure becomes valuable in the Great Lakes. Ports there are close substitutes and normally their aggregate capacity can absorb occasional demand shocks without wait times exploding. However, as extreme demand shocks become more frequent, all ports in the area can become congested. Since shippers have limited ability to substitute outside the region, investing in infrastructure becomes vital to alleviate congestion.

Related Literature This paper relates to a broad and diverse literature studying the role of transportation and infrastructure on trade; this includes work in international trade, port development, urban transport and congestion, as well as supply chain disruptions.

A recent and influential literature has focused on the impact of transportation infrastructure on

trade (e.g. Redding and Turner, 2015, Donaldson and Hornbeck, 2016, Redding, 2016, Donaldson, 2012, Fajgelbaum and Schaal, 2020, Allen and Arkolakis, 2022), as well as the evaluation of investment in infrastructure (Glaeser and Poterba 2021a and 2021b). Similarly to some of this work, we seek to quantify the welfare gains from investment in infrastructure, focusing on ports (instead of roads or railways). We contribute to this literature by explicitly modeling the dynamics of congestion and considering the role of macroeconomic volatility, while relying on a more granular methodology that captures the micro-structure of the infrastructure system.

A shorter literature has focused specifically on the impact of ports. Brooks, Gendron-Carrier, and Rua (2021) and Ducruet, Juhasz, Nagy, and Steinwender (2022) focus on the historical context of containerization to study the impact of port development on city population growth through higher market access, but also land and disamenity costs. Bailey (2021) considers competition between US ports focusing on business stealing. Similarly, Sytsma and Wilson (2020) consider the port choice of US importers. We extend the literature by incorporating a model for congestion in our analysis and allowing for positive spillovers across ports, while not taking a stance on how they compete. Fuchs and Wong (2022) explore investments in multi-modal transport networks relying on the methodology of Allen and Arkolakis (2022). See also Friedt and Wilson (2020) for a recent survey of the impact of port investment both on local economic activity, but also on the economy more broadly. A literature in port management and operations has studied related issues (Notteboom et al., 2022). Finally, we also contribute to the literature that studies the shipping industry, e.g. Hummels and Skiba (2004), Hummels (2007), Hummels, Lugovskyy, and Skiba (2009), Kalouptsidi (2014), Ishikawa and Tarui (2018), Asturias (2020), Brancaccio, Kalouptsidi, and Papageorgiou (2020), Heiland, Moxnes, Ulltveit-Moe, and Zi (2021), Ganapati, Wong, and Ziv (2022), Wong (2022) and Brancaccio et al. (2023).

In addition, we relate to a large literature that studies congestion, primarily in the context of urban transportation. In classic work, Vickrey (1969) and Arnott et al. (1990) provide models of road congestion, while more recently, a growing body of work embeds different types of congestion models in structural models of urban transportation (e.g. Durrmeyer and Martinez, 2022, Bordeu, 2023, Kreindler, 2023, Almagro et al., 2024). We rely on canonical models of queueing theory (see e.g. Larson and Odoni, 2007 and Bhat, 2008) and embed them in an empirical framework that allows us to analyze infrastructure investment.

We also relate to a nascent but growing literature on supply chain disruptions (Carvalho et al., 2021, Grossman et al., 2021, Alessandria et al., 2022, Bai et al., 2024). Finally, this paper naturally relates to

the large literature in international trade studying the role of geography in international trade, as well as the importance of trade costs in explaining trade flows between countries (e.g. Eaton and Kortum, 2002, Anderson and Van Wincoop, 2003, Melitz, 2003).

The rest of the paper is structured as follows: Section 2 briefly describes port operations, outlines our main data sources, and documents facts regarding time at port. Section 3 introduces our framework for port technology and port demand. Section 4 presents the estimation procedure and results. Sections 5 and 6 contain our main results on welfare gains from investment in infrastructure and net returns, while Section 7 concludes.

2 Industry, Data Description and Main Empirical Patterns

In Section 2.1 we describe the main aspects of port operations; in Section 2.2 we present the data sets used and some basic summary statistics; and in Section 2.3 we present facts about time at port.

2.1 Port Operations

A port is a maritime facility, consisting of both water and land, where ships can load or discharge cargo. Upon arrival at the port, a vessel may have to wait in anchorage, a designated area outside the port. Once berth space is freed up, the ship enters the port to get serviced: its cargo is loaded or discharged by port workers and port capital equipment, such as cranes. The cargo is stored in the port's designated storage areas to be later moved on to trucks or rail to its final destination, and the ship departs.

A port consists of a "landlord" (usually the local port authority), who owns the land and part (or all) of the basic port infrastructure, which may include berths, (un)loading equipment, storage areas, buildings, etc. Sometimes, the landlord leases land to terminal operators who are usually private enterprises. These are responsible for the numerous port services a ship receives while at port, such as pilotage, cargo-handling, towage, etc. A ship receives on average 40 services per visit, for which it pays different charges. For raw materials (bulk goods), which are the focus of this paper, port charges can be a substantial portion of their value, roughly equal to 15% (World Bank). Although they can differ substantially across ports, these charges do not vary much over time, as discussed below.[2]

Investment in port infrastructure is the focus of heated policy debates. Between 2012-2016, $46 billion were dedicated to capital investments (of which $18 billion was public), while between 2021-2025,

[2]There are a few explanations for this, related to political incentives to facilitate trade and maintain stability and transparency, as well as to the complexity of pricing, which involves a long list of services.

$163 billion expenditures are scheduled according to the American Association of Port Authorities. Port infrastructure involves water investments, which amount largely to dredging, as well as land investments, which relate to storage expansion.

In the US, port investment is usually funded by a mix of private and local public funds (e.g. through municipal bonds).[3] Investment decisions are undertaken by port (or other local) authorities and terminal operators who may have a mix of incentives including output or revenue considerations, job creation, local spillovers, and profit maximization. Port authorities face different oversight by state authorities; for instance, Oregon ports face substantial constraints by the state, as opposed to Florida ports where port authorities act largely independently.[4] Overall, in the US, ports should be thought of as quasi-public authorities, primarily aiming to serve local interests, such as high throughput, employment and connectivity to the area.

In this paper, we focus on commodities and raw materials, such as iron ore, steel, coal and grain, which are carried in dry bulk carriers. These are vessels designed to carry a homogeneous unpacked dry cargo, for individual shippers on non-scheduled routes. Bulk carriers operate like taxi cabs: each ship transports only one cargo at a time, for a trip between a single origin and a single destination.[5] The US is an important exporter of these goods and its ports handle a substantial amount of worldwide trade, especially grain which is exported worldwide from the ports in the Gulf. Overall, dry bulk shipping accounts for about half of total seaborne trade in tons (UNCTAD) and 45% of the total world fleet, which includes also containerships and oil tankers. The industry is unconcentrated, consisting of a large number of small shipowning firms (see Brancaccio et al. (2020) for a detailed description of the industry, as well as a conceptual framework for its operation).

Our analysis focuses on the largest 51 US ports by bulk throughput, covered by the datasets described below.[6] Ports are very diverse along several dimensions. Geography is a key determinant of a port's performance: its location and proximity to producing and consuming regions, but also whether it is in deep vs. shallow waters, inland vs. sea, its tide range, etc. The left panel of Figure 1, depicts the annual

[3] Waterways are an exception, as they are federally funded.

[4] https://www.oregon.gov/biz/Publications/Ports/2010PortPlan.pdf

[5] It is worth noting that bulk ships are very different from containerships, which operate like buses: containerships carry cargos (mostly manufactured goods) from many different cargo owners in container boxes, along fixed itineraries/routes according to a timetable. It is technologically impossible to substitute bulk with container shipping. We choose to focus on bulk carriers because they offer a simpler setup to study port congestion (as would oil tankers). Indeed, container shipping features many complications such as hub-and-spoke networks of ports, vertical integration with port operators, market power against ports and others. Nonetheless the basic economic mechanisms discussed in this paper are useful for containers as well, as well as other settings with congestion, such as airports.

[6] We observe port calls for all 61 US ports that handle bulk. However, for 10 of these ports (which account for 5.4% of total ship arrivals) we were not able to obtain data on labor or infrastructure, so we drop them from our analysis.

throughput of US ports and illustrates the extent of size heterogeneity: the smallest ports in this group handle about 10 ship per year, while the largest handle more than 600 per year.

2.2 Data

We collect a number of novel datasets that are necessary to estimate the three key ingredients of our analysis: the (i) port production function; (ii) port demand; (iii) cost of port infrastructure.

Data for Port Production Function In order to measure how long ships spend at the port, we employ the universe of port calls, between 2016 and 2021 for bulk carriers larger than 10,000 DWT. The dataset is obtained from AXS Marine (AXS Dry). Each observation concerns a unique vessel and port, and reports a timestamp for the vessel's arrival at the port's anchorage, a timestamp for the vessel's entry to and the exit from the port, as well as an indicator for whether the ship loaded or discharged, and an estimate for the commodity on board. The port calls are created by AXS Marine from AIS (Automatic Identification System) data that report exact vessel positions (latitude and longitude) every few minutes.

In addition, we construct a dataset of global port infrastructure.[7] Because ports seldom share information about their infrastructure, we use geolocation data and satellite imagery (provided by Google Maps and Google Earth Pro) to collect this information. Focusing on the US, we collect data for each of the ports in our sample, from 2010 through 2021, and manually encode the following measures of port infrastructure: miles of each berth, acres of storage space, and number of cranes. In Appendix B we provide the details on how this dataset was collected and encoded.

Finally, to obtain a measure of port employment, we use publicly available annual data on longshoremen union membership reported to the Department of Labor in the US. We match this to ports based on the zipcode of each union.[8]

Data for Port Demand We leverage the universe of port calls to construct the monthly number of arrivals at each of the ports in our dataset. Because this dataset does not contain information on the

[7] Although there are publicly available datasets that report some measures of port infrastructure we concluded that the reported measures did not accurately reflect a port's capacity. To give one example, a commonly reported capacity measure is the number of berths; berths however can be short or long. For example, berths at the port of Baltimore (MD) range from 0.12 miles to 1.0 miles, and clearly offer different capacity. In addition, berths are not the only relevant feature of a port's capacity: storage space is also important, since limited storage space implies a ship may need to wait until storage is emptied out (e.g. by trucks) to further discharge. Cranes are another important measure of port infrastructure.

[8] This dataset is not as granular as we would like, since membership is reported annually, and it only includes longshoremen that are fully registered (non-casual) union members. We validate our findings using detailed employment data that is available for four large US ports (Oakland, Los Angeles, Tacoma and Seattle) by the Pacific Maritime Association. Indeed, as shown in Table 4 in Appendix A, total man-hours from shift-level information is very highly correlated with our annual data.

final destination (or origin) of a shipment within the US, but only the US port used (and the foreign location from where the ship came from), we complement it with the Freight Analysis Framework (FAF) Regional Database of the Bureau of Transportation Statistics, which contains detailed data on freight movements between foreign regions and US states and metropolitan areas, including information on the region of entry (or exit) as well as the product transported.[9]

Additionally, we use a dataset of port charges, from a provider that brokers such transactions and who wishes to remain anonymous. This is a sample of port calls globally that covers approximately 35% of port calls, with the associated breakdown of the charges for each service requested (of which the main ones are berth dues, which include cargo handling dues, line handling, and pilotage). The dataset runs from 2016 to 2020. Although we cannot guarantee that this sample is not selected, we find that it has similar geographical and ship size coverage as our universal port call data.

Data for Cost of Port Infrastructure Port expansion amounts to constructing an additional section of berth to service incoming or outgoing freight. The cost of such expansion consists of two main elements: (i) dredging; and (ii) acquisition of land for storage. For the former, we use publicly available cost projections from US Army Corp of Engineers (USACE) port projects. From these, we produce site-specific cost estimates for dredging a new berth, taking into account the rock vs. non-rock content of each port-berth's seafloor, obtained from the US Geological Service's (USGS) National Crustal Model. For the cost of land acquisition, we first use open-source map data from OpenStreetMap, to identify parcels of industrial and commercial land near each existing berth from our hand-collected geolocational data from Google Earth. We then use highly granular land value estimates from the ecology and conservation literature (Nolte, 2020) to assign unit land prices to these parcels.[10] We discuss our approach in detail in Section 6.1, as well as Appendix D.

Summary statistics Our final port call dataset involves 42,341 port calls between 2016 and 2021, for 7,499 ships and 51 ports. As shown in Table 1, on average, ports handle 136 separate port calls (and about 7.3 thousand tons) annually, though this can vary substantially (as shown also in the left panel of Figure 1). Ships spend on average 117 hours at port; of this, on average 67% is attributed to service time and 33% to wait time at anchorage (i.e. time spent in congestion). However, time at port

[9] FAF groups foreign countries in eight international regions: Canada, Mexico, Rest of Americas, Europe, Africa, South-West & Central Asia, Eastern Asia, and South East Asia & Oceania.

[10] A similar approach was taken by Ducruet et al. (2022) who show the criticality of land scarcity and assumed the port expansion cost amounts to land costs. We extend their cost analysis by adding dredging costs, which amount to 26% of total costs based on our estimates.

	Mean	SD	25th percentile	75th percentile
Annual number of port calls	136	130	54	180
Annual throughput (tons)	7,295,140	7,355,597	2,602,124	9,970,174
Time at port (hours): average	117	54	74	153
Time at port (hours): st. dev. over time	88	41	58	118
Wait time (hours): average	39	39	2.9	58
Wait time (hours): st. dev. over time	58	50	16	101
Fraction of port calls with positive wait time	41%	30%	13%	70%
Port dues ($ per ton)	2.1	1.3	1.2	2.9
Port dues (in $) per port call	117,345	96,206	44,038	168,399
Length of berths (miles)	0.96	0.8	0.42	1.4
Storage space (acres)	176	176	55	234

Table 1: Summary Statistics. Port calls (throughput in tons) is the total annual number of ships (total tonnage) that were handled at port. The average time at port refers to time at port averaged over time, while standard deviation of time at port is computed over time per port. The fraction of port calls with positive wait time is computed for each port and year.

varies dramatically across ports and years. Finally, it is also worth noting that wait time at anchorage is substantially more volatile than service time at the port level.

Total port charges are on average $117,345 per port call. Port prices are heterogeneous over ports, but not over time. Indeed, in Table 3 in Appendix A we regress port prices on several observables (e.g. port, month, commodity, ship size) and demonstrate that port fixed effects soak much of the price variation, in contrast to time fixed effects that tend to explain little (similarly, there is no evidence of additional charges during congested periods).

On average, ports have 0.96 miles of berths (sd 0.8) and 176 acres of storage (sd 176), however the heterogeneity in infrastructure mirrors that in size; e.g., the port of Charleston (SC) has 27 acres of storage and 0.14 miles of berths, while the port of Houston (TX), one of the largest in the US, has 290 acres of storage and 1.9 miles of berths.

Finally, ports do not tend to specialize in specific commodities; 77% of ports receive all major bulk commodities including grain, ore, and coal (note that containers are handled in different terminals).[11]

[11]Using the granular data on port operations available for West Coast ports from PMA, we test whether bulk operations are independent of container operations within a port; see Table 6 in Appendix A. In particular, we find that the number of container ships serviced in a port in a given shift does not affect the service time for bulk ships serviced in the same shift. This suggests that bulk and container operations are handled at different terminals, possibly due to the different technology necessary to service these different types of cargoes.

2.3 Time at Port

Our key metric of port efficiency is the total time a ship spends at a port, defined as the sum of the wait time, i.e. the time a ship spends waiting in anchorage and the service time, i.e. the time a ship spends inside the port discharging or loading its cargo. Time at port is naturally a key determinant of port throughput.[12]

Time at port is extremely costly to ships and cargo-owners. The 117 hours spent on average at US ports constitutes 22% of travel time, suggesting that iceberg transport costs that are a function of distance miss a key component of trade costs (almost half -roughly 44%- if both loading and unloading ports are taken into account). At the average shipping price of $14,000 per day during our sample, time at port has a direct cost of $68,250 per port call, plus any penalty fees ("demurrage") paid to/from shipowners in case of delay. Naturally, this is not but a lower bound of the costs accrued, since it does not include delays for the next intermodal steps (usually trucks or rail), nor inventory costs that importers and exporters face. The demand estimates in Section 4.2 allow us to quantify directly the cost of these disruptions.

Although informative, the average time at port masks massive dispersion over both space and time. The left panel of Figure 1, depicts the average time at port in the US and demonstrates the sizable differences across ports: annual average time at port can range from 7 to 240 hours (see also Table 1). For the average (median) port 41% (46%) of the ships that arrive face a queue; however in some ports, such as Corpus Christi (TX), ships face queues more than 67% of the time, while others, such as Cleveland (OH), they rarely wait in anchorage, if ever. This large heterogeneity is reflected in the different returns from investing in infrastructure, discussed in Section 5.

Similarly, time at port exhibits significant fluctuations over time, as shown in the right panel of Figure 1. Perhaps not surprisingly, the largest uptick in time at port occurs during the COVID-19 pandemic whereby ports experienced substantial increases in congestion. Indeed, the fraction of ships that had to wait almost doubled, going from 36% in 2019 to 77% in 2021. It is worth noting that these graphs do not capture the large rise in congestion observed in 2022 which is outside of our data sample. Finally, a key feature of the port network is that ports tend to get congested at the same time: as an example, for Baltimore (MD) and Hampton Roads (VA), two of the largest ports in the East Coast, the correlation of wait time is 0.48. This pattern suggests that macro volatility is crucial and generates correlated shocks across ports.

[12] We focus on time at port rather than port throughput because the former contains more information than the latter. To see this, consider two ports that handle five ships in a month, but in one port, time at port per ship is four days, while in the other it is ten.

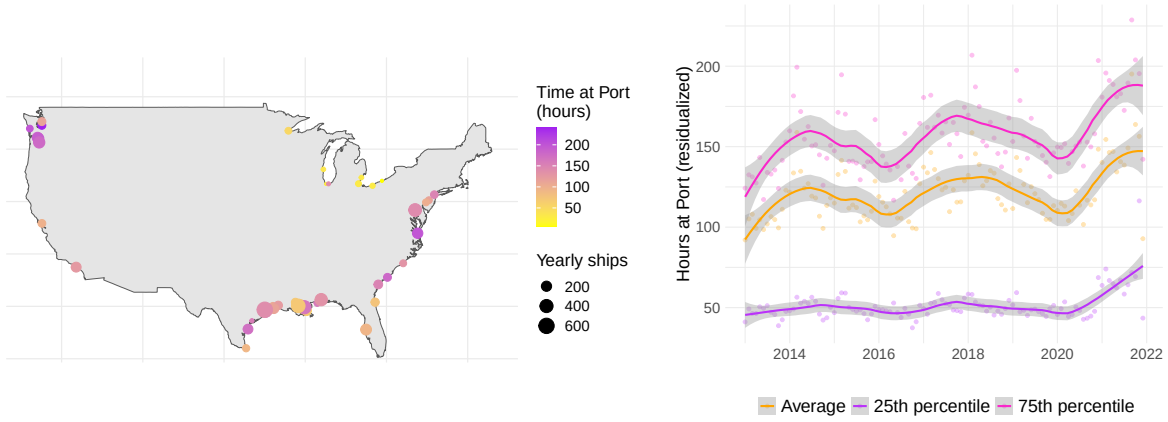

Figure 1: Port throughput and time at port over space and time. The left panel plots the average time at port in US ports. It also indicates the size of the port in terms of annual throughput, designated by the size of its dot. The right panel plots the average time at port for each month, as well as its 25th and 75th percentile. We control for ship size and commodity on board.

3 Framework

In this section we present our basic framework. Importer-exporter pairs decide whether to trade between a foreign and a US domestic location, and which US port to use. This choice depends on the geographical location of the ports, congestion there, as well as price and port quality. When the ship arrives at port, it either enters the anchorage and joins the queue, or it enters directly for service, if there is no queue. Ships that get serviced depart. The port queue evolves endogenously as ships arrive at port or depart after service. We first describe port technology, and then the demand that ports face.

3.1 Port Technology

Port technology transforms port inputs, such as labor, cranes and infrastructure, to port output, i.e., time at port. Consider a ship arriving at port. As already discussed, its time at the port is given by the sum of the service time, i.e, the time it takes for the ship to be loaded or discharged, and the wait time, i.e. any time the ship spends in anchorage, waiting to begin service.

We formalize this by adopting one of the most widely-used queueing models, the so-called the M/M/K model, where ship arrivals follow a Poisson process and port service times are exponentially distributed (see e.g. Larson and Odoni, 2007 and Bhat, 2008). Suppose that a ship arrives at port and finds Q ships ahead, either waiting in anchorage or at berth, each with expected service time T, while the port can

handle at most K ships at a time. Then the ship's expected total time at port is equal to,

$$\underbrace{T}_{\text{service time}} + \mathbb{I}\{Q \geq K\} \underbrace{(Q - K + 1)\frac{T}{K}}_{\text{wait time (queueing)}} \tag{1}$$

In words, if the ship arrives and there are fewer than K ships at port (i.e. $Q < K$), then the port is operating under capacity, and the ship is able to be serviced immediately, so that total time at port simply equals the service time, T. If instead $Q \geq K$, then the ship needs to wait for a berth to open up and the total time is augmented by the wait time.

To understand the expression for wait time in equation (1), note that if the port could handle only one ship at a time ($K = 1$), then the expected wait time would just equal QT, as the ship would need to wait for Q ships to complete service, each taking on average T periods. In the case where $K > 1$, the arriving ship does not have to wait for all Q ships ahead to complete service, since the port can handle more than one ship at a time. For example, if $Q = K$, then only one ship needs to complete service for a berth to open up, regardless of how large Q is. In general, $Q - K + 1$ ships need to conclude service before the ship can get to berth.[13] Since service times are exponentially distributed, a berth becomes available in T/K periods in expectation when K ships are being serviced. Therefore the expected wait time is $(Q - K + 1)(T/K)$.[14]

Expected service time at port is given by

$$T(L/s, c/s, \omega) \tag{2}$$

and depends on the port's workers, L, cranes c, productivity ω, as well as the number $s \leq K$ of ships currently serviced. This specification is driven by the empirical finding that a ship's service time depends on how many other ships are being serviced (see Table 5 in Appendix A), suggesting that ports to some extent stretch their inputs across ships handled.[15]

[13] To see this, suppose that $K = 3$ and there are $Q = 7$ ships ahead, so that three ships are currently being serviced and there are four ships ahead in anchorage. Then, the three ships currently at berth need to conclude service, and be replaced by three of the ships in anchorage, leaving only one ship ahead in anchorage. A berth now becomes available for the arriving ship, once two of the ships at berth conclude service. Therefore, only five ships need to complete service before a berth becomes available.

[14] In the M/M/K queueing model a queue forms in a system with K servers; inter-arrival times and service times are exponential. Equation (1) relies on the memoryless property of the exponential distribution (see Bhat, 2008).

[15] There are some implicit assumptions behind our formulations in (1) and (2), all of which have been verified empirically. In particular, we do not allow service time, T, to be affected by the number of ships in line (see Table 5 in Appendix A). Moreover we do not allow ships to skip the line so that the line is first-in-first-out (FIFO), as we find that in the US queue jumping occurs for less than 20% of observations. Finally, we do not allow port capacity K to adjust depending on the number of ships waiting.

In summary, time at port has three key components: service time T, port capacity, K, and Q which is endogenous and depends on T, K, as well as the demand for port services. In our empirical implementation, we allow each of these objects to vary across time and ports.[16] From equations (1) and (2) the differential impact of different inputs starts to appear: infrastructure in port capacity, K, matters for the port queue and therefore waiting time, whereas workers and cranes primarily affect the service time of each ship.

Finally, we discuss how the port queue evolves endogenously as ships arrive, are serviced, and leave. To this end, denote by λ the ship arrival rate and by $\mu(s) = 1/T(s)$, the rate at which a ship completes service, where we suppress the dependence of $T(s)$ on L, c, ω.[17] Then, the evolution of the queue follows a Markov structure. Indeed the probability $p_n(\tau)$ that in instant τ there are n total ships at port, either being serviced or waiting, satisfies,

$$\frac{dp_n(\tau)}{d\tau} = \begin{cases} -\lambda p_0(\tau) + \mu(0)\, p_1(\tau), & \text{if } n = 0 \\ -(\lambda + n\mu(n))\, p_n(\tau) + \lambda p_{n-1}(\tau) + (n+1)\,\mu(n) p_{n+1}(\tau), & \text{if } 0 < n < K \\ -(\lambda + K\mu(K))\, p_n(\tau) + \lambda p_{n-1}(\tau) + K\mu(K) p_{n+1}(\tau), & \text{if } n \geq K \end{cases} \tag{3}$$

In words, consider the case where $n < K$ (a similar argument hold for $n \geq K$), then the change in the probability that there are n ships at port increases by (i) the probability that there are $n-1$ ships at port at τ, $p_{n-1}(\tau)$, and a new ship arrives, which happens with probability λ; and (ii) the probability that there are $n+1$ ships at τ, $p_{n+1}(\tau)$, and one of them is serviced, which happens with probability $(n+1)\,\mu(n)$. In addition the probability $p_n(\tau)$ decreases by the probability that there are n ships at port, $p_n(\tau)$, and a new ship arrives or one is serviced. Therefore, simulating the queue evolution boils down to solving this system of differential equations. We describe these simulations in detail in Appendix C.[18]

[16] Throughout the model, we abstract away from ship size. Indeed, the number of "servers", K, would not be sufficient to describe the port's capacity, and one would have to consider the different combinations of ship sizes that could "fit" in the port at a single point in time. At the same time, each ship would need to keep track of the distribution of ships ahead of it, not only their number. This is a complex problem that even modern ports and airports do not necessarily perfectly optimize over. At the same time, capturing ship size does not impose first-order limitations to our analysis. Indeed, the focus of the paper is to compute welfare gains of expanding a port's capacity, holding fixed the distribution of ship sizes that a port handles.

[17] We suppress the dependence of μ and T on L, c, ω, to ease exposition, as we assume that these variables change at a much lower frequency than the queue at the port which changes with each ship arrival. We also suppress the dependence of the ship arrival rate, λ on time, though in practice we allow demand to vary over time as we discuss in Section 3.2 (see also discussion in section 4.11 of Larson and Odoni, 2007). In Appendix C we describe the queue simulation in detail. Note also that the dependence of μ on s means that we use a general birth and death queuing model, which nests the case of M/M/K.

[18] We require that there is a high enough queue length $\bar{L}$, above which ships choose not to join the queue and $\lambda = 0$ in this

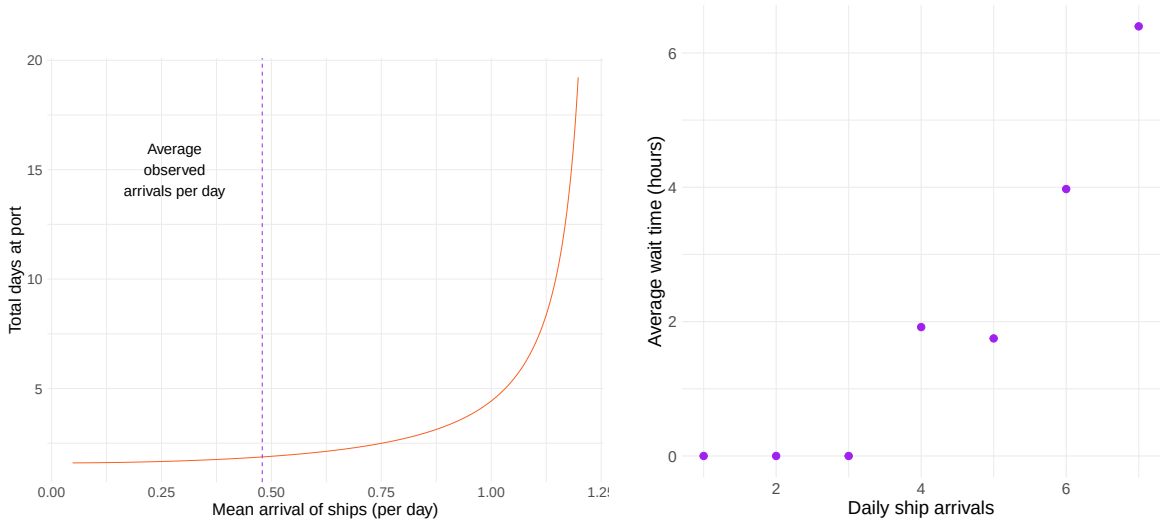

Figure 2: Convex cost of congestion. The left panel of this graph plots the relationship between time at port (vertical axis) and the mean arrival rate (horizontal axis) and is produced from model simulations of increasing the mean arrival rate of ships. In particular, for each level of the mean arrival rate on the horizontal axis, we simulate the port queue using the system of differential equations (3). We use the average service time, and the average capacity K. The right panel plots the empirical counterpart of this convex cost of congestion in the port of Duluth (MN). In particular, it plots the average time ships spend waiting (vertical axis) against a moving average of daily ship arrivals (horizontal axis).

3.1.1 Understanding Congestion Dynamics

This framework sheds light on the formation and evolution of congestion, which we demonstrate through a series of simple simulation exercises. These insights will prove useful in understanding the impact of investment in infrastructure.

Convex Congestion Cost Function First, consider the relationship between port congestion and demand. The left panel of Figure 2 plots the simulated time at port for the average port, for different values of the ship arrival rate.[19] Since time at port is costly, this graph captures the cost of congestion.

In periods with low average demand, time at port is fairly low: as demand is low, it is likely that the port is not at capacity, so an arriving ship immediately begins service and total time at port simply equals the service time. Note that increases in demand do not initially change time at port much, since the probability that the port is full remains low. After a certain point, however, it becomes more and more likely that an arriving ship has to wait at anchorage and total time at port increases, as the wait time

case. Thus,

$$\frac{dp_{L+K}(\tau)}{d\tau} = -K\mu(K)\, p_{L+K}(\tau) + \lambda p_{L+K-1}(\tau)$$

In practice, we choose $\bar{L}$ to be very large so that it is never binding.

[19] We construct the average port as a port with average capacity, ship arrival rate, and service time. For each level of the mean arrival rate, we simulate the queue using equations (3).

becomes positive. Further increases in demand have large effects on wait time, as the port infrastructure becomes overwhelmed and the queue explodes.[20] In other words, even if we assume that for customers the cost of waiting at port is linear in time at port, the port production function endogenously generates a convex cost function for congestion.

This convex congestion cost function explains why ports are susceptible to disruptions: at low levels of port utilization, demand shocks have little impact on the total time at port; on the contrary, when the port is already at capacity, further increases in demand can have large effects on the queue and ships' wait time. Therefore, an implication of this model is that even small demand shocks can lead to significant disruptions, if they occur when the port is already close to capacity. By the same logic, the accumulation of shocks can prove detrimental for port congestion.

This model formulation provides a microfoundation for the congestion costs encountered in the literature. For instance, in road congestion models, congestion often depends parametrically on the current demand for road usage. A notable difference is that here the cost of congestion depends not only on current demand for port services, but also on past demand realizations, allowing for endogenous congestion dynamics.

Although this analysis is based on model simulations, we find that this convexity of congestion holds in the data as well. We depict one example in the right panel of Figure 2, which shows this relationship for the port of Duluth (MN). The figure plots a moving average of daily ship arrivals on the horizontal axis and the average time ships spend waiting on the vertical axis. As predicted by our framework and the simulations, in periods of low demand, ship wait time is zero, as the port has spare capacity and ships immediately begin service; as demand increases, ships are more and more likely to have to wait. Interestingly, as demand increases further, wait time seems to explode, resulting in the expected convex relationship between time at port and port congestion.

The Role of Variance Next we examine the impact of an increase in demand volatility. Formally, we compute the average monthly arrival rate across ports and simulate wait times under a mean preserving spread for this process, where the dispersion of these monthly arrival rates increases. This implies that both very low and very high demand realizations now become more likely. Note that, due to the convexity of time at port with respect to ships' arrivals, high and low demand realizations have an asymmetric impact on time at port. On one hand, time at port is relatively flat for low realizations of ships' arrivals; on the

[20] Formally the ratio of ships entering the system over the ships exiting approaches one, in which case the port queue approaches infinity.

other hand, high demand shocks dramatically increase congestion, overwhelming the system. As a result, increasing the variance of demand increases the average time at port that ships face. This is illustrated in the left panel of Figure 3 and suggests that it is not just the level of demand, but also its variance that increases the cost of congestion.

Convexity and Infrastructure How does investment in infrastructure impact congestion costs? Consider an improvement in infrastructure, i.e. an increase in port capacity K. As shown in the right panel of Figure 3, the impact of this improvement is highly asymmetric for low vs. high demand. Indeed, higher capacity leads to very little change in time at port in states of low demand: time at port is roughly equal to the service time, which remains unaffected when K increases. When demand is high, on the other hand, there can be a substantial impact on time at port, as wait times fall. Much like a super market queue would disperse if a new cashier opens up, port wait time is slashed when new infrastructure comes online. Capacity, K, essentially only matters in periods where the port is already constrained, acting as a form of insurance against periods of high demand realizations, that can smooth out volatility.

It is worth contrasting the impact of infrastructure to that of a decline in service time T, which from (2) can be achieved by an increase in some combination of labor L, cranes, c and/or port productivity, ω. A reduction in service time is beneficial at all levels of demand, as shown in the right panel of Figure 3. At low levels of demand where total time at port is largely comprised of the service time, there is an immediate reduction in total time. At higher levels of demand, total time at port falls as well: beyond own service time being lower, the service time of ships ahead in the queue also falls, thus reducing wait time as well, leading to a substantial reduction in time at port. Hence, reductions in service time, generated from higher variable inputs, lead to symmetric reductions of congestion costs, in contrast to the case of improvements in infrastructure K.

3.2 Demand for Ports

We next turn to demand for port services, which is necessary to estimate the welfare gains from port improvements.

There are D locations in the US, including both coastal and inland regions, and F locations abroad. For every pair (d, f) of US (domestic) and foreign locations, every month there are M_{df} importer-exporter pairs that have the option of trading a ship-sized shipment between the two locations, either originating in d, or destined there. In the case of exports, d is a US producing location and f is a foreign destination,

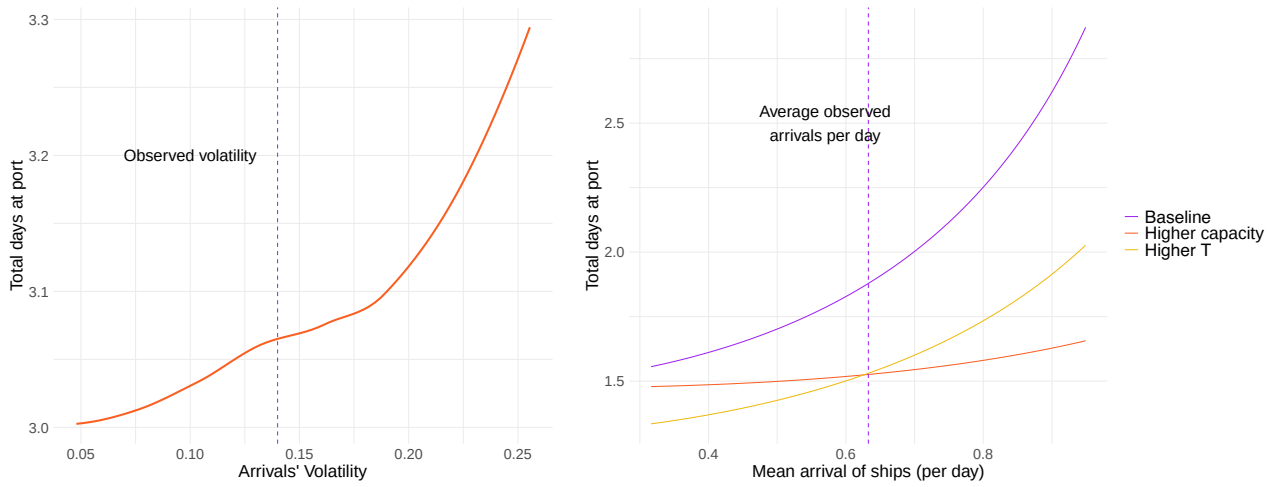

Figure 3: Congestion dynamics: the role of volatility and different port inputs. The left panel plots the relationship between time at port (vertical axis) and the variance of the arrival rate (horizontal axis). It is produced from model simulations of increasing mean preserving spreads of the arrival rate of ships. The right panel plots the relationship between time at port (vertical axis) and mean arrival rate (horizontal axis) when service time T falls or capacity K increases.

while in the case of imports, d is a US importing destination and f is a foreign producing location. Let t denote the month-year period. Each shipment i needs to decide whether to trade and, if so, which port, $j \in \{1, ..., J\}$ to use. The payoff to shipment i, traveling between d and f, from choosing port j in period t is given by,

$$u_{ijt} = \beta_f \text{dist}(f, j) + \beta_d \text{dist}(j, d) - \beta_T TT_{jt} - \beta_p p_{jt} + \gamma_t + \gamma_f + \gamma_{l(j)} + \xi_{jft} + \epsilon_{ijt} \tag{4}$$

where $\text{dist}(f, j)$ denotes the distance between the foreign location f and port j and$\text{dist}(j, d)$ is analogously defined, TT_{jt} is the total time at port j in period t, and p_{jt} is the price that port j charges for its services in period t. The β parameters capture how much the pair values each characteristic. Moreover, γ_t is a month fixed effect capturing macroeconomic fluctuations, while γ_f and $\gamma_{l(j)}$ are respectively fixed effects for the foreign location and the port's region (namely, either the Gulf, the East Coast, the West Coast, or the Great Lakes). Finally, ξ_{jft} is an unobserved demand shock for period t and ϵ_{ijt} is an iid shipment-specific shock that follows a Type I extreme value distribution. If the shipper chooses not to trade to the US, they obtain a payoff that is normalized to zero, so that $u_{i0t} = 0 + \epsilon_{i0t}$.

Note that TT_{jt} is endogenous since it depends on port demand, and therefore determined in equilibrium. In particular, it depends on the decisions of all other shippers; in equilibrium every shipper can correctly predict the port flows and infer TT_{jt}.

The specification in (4) allows us to measure the welfare cost of delays, since the term $\beta_T TT_{jt}$, which

captures the cost of port disruptions for the importer-exporter pair, also reflects the cost of delays along the entire supply chain. In addition, our specification allows for a rich demand response to an increase in time at port: a shipment's foreign and domestic locations shape how substitutable port j is when hit by disruptions. For example, when a shipment is headed to a final destination near port j, an exporter may be more willing to endure the long wait times rather than switch to a faraway port. Conversely, the shipper might be more willing to switch ports when the final destination is in the middle of the country.

Trade flows The payoff specification (4) implies that the probability that a pair trading between foreign location f and domestic location d chooses port j is,

$$\sigma_{jfdt} \equiv \Pr(j|f,d,t) = \frac{\exp(\delta_{jft} + \beta_d \text{dist}(j,d))}{1 + \sum_l \exp(\delta_{lft} + \beta_d \text{dist}(l,d))} \tag{5}$$

where following Berry et al. (1995) we define,

$$\delta_{jft} \equiv \beta_f \text{dist}(f,j) - \beta_T TT_{jt} - \beta_p p_{jt} + \gamma_t + \gamma_f + \gamma_{l(j)} + \xi_{jft}$$

which captures the "mean utility" of port j for shipments from/to f in month t. The probability of the outside option, σ_{0fdt}, is similarly defined. This implies that every month the total flow through j is given by $\sum_f \sum_d \sigma_{jfdt} M_{df}$.

Looking ahead at the estimation, one complication of this framework is that our data on ship flows, do not contain information on a shipment's US location d, so we do not directly observe $\Pr(j|f,d,t)$. In other words, we might know that a ship began from China and unloaded at the port of Long Beach (CA), but we do not know the shipment's final destination within the US; similarly in the case of exports, we know that a shipment loaded in Houston (TX) and traveled to Europe, but we do not know where in the US it was originated.

To overcome this challenge, we treat the domestic location d as a latent variable that needs to be estimated (this is reminiscent of unobserved consumer types as in Berry et al., 2006). In particular, marginalizing (5) by the probability, $\Pr(d|f,t)$, of a shipment to/from any domestic US location d given the foreign location f, we obtain,

$$\sigma_{jft} \equiv \Pr(j|f,t) = \sum_d \Pr(j|f,d,t) \Pr(d|f,t) \tag{6}$$

Note that $\Pr(j|f,t)$ is observed in our ship flow data. In order to measure $\Pr(d|f,t)$, we leverage the FAF Regional Database, which contains information on freight movements among major US regions. In particular, the FAF reports yearly trade flows between US regions, such as states or metropolitan areas, as well as flows between these regions and broad foreign locations. For instance, it reports flows from California to Southeast Asia or Europe, which allows us to measure $\Pr(d|f,t)$. These moments will be useful in writing down our model's likelihood as we discuss in Section 4.2 below.

Finally, the specific timing of the $\sum_f \sum_d \sigma_{jfdt} M_{df}$ ships that pass through port j within month t is random. In particular, a ship arrives at port j at Poisson rate $\lambda_{jt} = \sum_f \sum_d \sigma_{jfdt} M_{df}$.

Adding macroeconomic volatility Macroeconomic fluctuations greatly affect overall demand for ports and are thus central to the decision to invest in port infrastructure (recall Figure 3). For this reason, we enrich our framework in order to capture the evolution of aggregate demand. In our model, macroeconomic conditions are captured by the term γ_t. This term changes over time and is common across all ports j; intuitively, therefore, γ_t determines the choice of whether to export or not.[21] To capture the evolution of these macro shocks we assume that

$$\gamma_t = \rho_0 + \rho_1 \gamma_{t-1} + \epsilon_t, \tag{7}$$

with $\epsilon_t \sim N(0, \sigma_\epsilon^2)$. In Section 4.1 we describe how we augment our data to estimate the times-series process for γ_t.

4 Estimation and Results

We next discuss how we take our framework to the data, as well as the results from our estimation procedure.

4.1 Port Technology Estimation and Results

Service Time We specify port service time as a general CES specification where

$$T_{jt} = \omega_{jt} \left(\alpha \left(\frac{L_{jt}}{s_{jt}} \right)^{\eta} + (1-\alpha) \left(\frac{c_{jt}}{s_{jt}} \right)^{\eta} \right)^{-\frac{1}{\eta}} \tag{8}$$

[21] In fact, it would have been equivalent to define the payoff of the outside option as $u_{i0t} = \gamma_t + \epsilon_{i0t}$. It is however easier in the estimation to subtract γ_t from the utility of all alternative options j, and set the mean payoff from the outside option equal to zero, so that $u_{i0t} = \epsilon_{i0t}$.

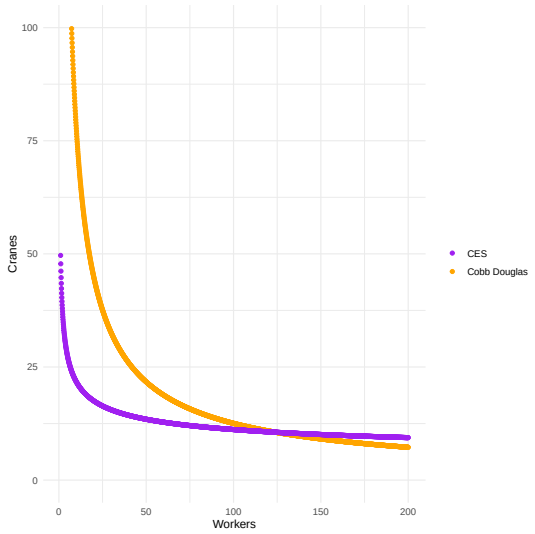

Figure 4: Port service time. The figure plots the isoquants associated with the estimated service time CES specification (8) (for the average service time); it also plots the corresponding isoquant if we were to estimate a Cobb-Douglas specification (i.e. $\eta \to 0$).

This encompasses both Cobb-Douglas ($\eta \to 0$) and Leontief ($\eta \to -\infty$) as special cases. Furthermore, following Arellano and Bond (1991), we let port productivity satisfy, $\log \omega_{jt} = \log \omega_j + \delta t + \nu_{jt}$, where ν_{jt} is iid with zero mean.

Following the usual concern when estimating production functions that inputs are not orthogonal to realized productivity, we introduce instruments, z_{jt}. First, we use local wages in related occupations as an instrument for labor. In particular, wage changes in occupations such as truck drivers, construction equipment operators, and taxi drivers, the most common substitute occupations for longshoremen (as indicated by occupational transition matrices), change their outside options and thus affect port employment. We construct this instrument using wage data from the Occupational Employment and Wage Statistics of the BLS and occupational switching probability data from the Current Population Survey (CPS) of the US Census Bureau. As additional shifters for a port's labor force, we use changes to the labor force and to the fraction of the population above 65 years old in the port county. Finally, we include lagged inputs as instruments. We leverage these instruments to estimate equation (8) using a GMM estimator based on the moments $\mathbb{E}(\nu_{jt} z_{jt}) = 0$.[22]

As shown in Figure 4, our results suggest that labor and cranes have low substitutability, as the shape of the isoquant is more similar to that of a Leontief (the estimates are in Table 8 in Appendix A).[23]

[22] Since we have annual data on labor and cranes, and our instruments are also annual, we estimate equation (8) at the port-year level. In particular, we residualize, separately for each port, the observed service times on ship size and commodity carried and use its yearly average for estimation.

[23] As is typical in most of the literature, we do not allow the shape of the production function to vary over time or across units. Using detailed data on four large US ports (Oakland, Los Angeles, Tacoma and Seattle) from the Pacific Maritime

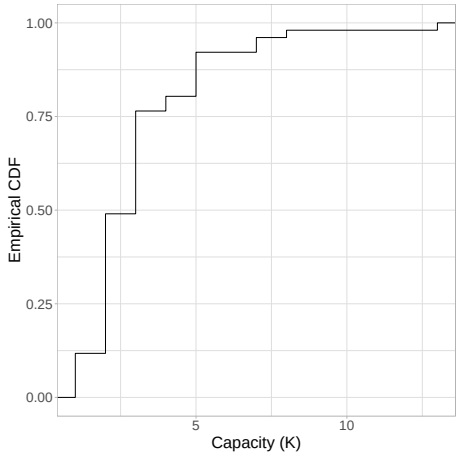

Figure 5: Port Capacity. This figure plots the empirical cumulative distribution of ports' effective capacity K_j, measured using AIS data, as the maximum number of ships that can be handled simultaneously, conditional on at least one ship waiting in anchorage.

Moreover, we find evidence of substantial productivity dispersion: as shown in Figure 7 in Appendix A, there is sizable heterogeneity across ports in the US, even within the same region. The same figure shows that our port productivity estimates are consistent with anecdotal evidence that ports that have lower tide range are more productive. Moreover, ports on rivers and ports with man-made breakwater structures are less productive than natural coastal ports.

Port Capacity In our setup, effective port capacity, K_j, is the maximum number of ships that a port can handle at a time. Due to our granular port call data, we can measure K_j directly, instead of using other common but poor proxies, such as the number of berths. For every port we consider the number of ships simultaneously being serviced, conditional on at least one ship waiting outside the port. We condition on ships waiting, since this implies there is no space inside the port and therefore we avoid underestimating port capacity. We then take the maximum of this measure over the sample period and this pins down K_j.

Figure 5 plots the empirical distribution of recovered capacity across ports. Most ports have low capacity: 75% of ports can service at most 3 ships simultaneously. The largest ports, however, have substantially higher capacity; for instance Houston (TX) and Baltimore (MD) can handle up to 8 and 13 ships respectively.

The port's effective capacity K_j is a function of its physical infrastructure, such as the length of its berths and storage space. Estimating the capacity function is crucial when we later gauge the cost of

Association, we find supportive evidence for this assumption: the share of crane operators over total labor is approximately 15% and varies little both over time and across these ports. This suggests that there is not much heterogeneity over the technological relationship between workers and cranes across ports. For more details see Figure 17 in Appendix A.

increasing K_j, derived from dredging and acquiring land.

We assume that capacity is given by a Leontief function: a ship requires both space to park (length of berth), as well as the land space to (un)load the cargo, and it is unlikely that the two are flexibly substitutable. Therefore we have,[24]

$$K_j = \min\left\{\frac{B_j}{\kappa_1}, \frac{A_j}{\kappa_2}\right\} \tag{9}$$

where B_j is port j's total length of its berths, A_j is its acreage of storage space, and κ_1, κ_2 are parameters.[25] In particular, κ_1 and κ_2 are technological constants determining, respectively, the mileage of berths and acreage of storage required to handle a ship.

An implicit assumption here is that the parameters κ_1, κ_2 are the same across ports; this is consistent with data: Figure 19 in Appendix A suggests that the length of berths per ship, as well as the storage space per ship are roughly similar across ports. With that in mind, we use our infrastructure data to calibrate these parameters. We set κ_1 equal to the median number of miles of berth per ship, B_j/K_j, and κ_2 equal to the median storage acreage per ship, A_j/K_j.[26]

The Leontief specification proves a very good fit. To showcase this, we also estimate a CES specification (which nests Leontief) for the capacity function. As shown in Figure 18 in Appendix A, which plots the isoquants for $K_j = 1$, the CES ends up being very close to a Leontief. This is consistent with our intuition that length of berth and storage space are not substitutable and are needed in fixed proportions.

4.2 Demand Estimation and Results

Based on equation (6), we formulate a likelihood for the port call microdata as follows:[27]

$$\begin{aligned}
\mathcal{L}^{\text{micro}}(\delta, \beta_d) &= \sum_{f,t}\sum_{j} n_{ft}^{j} \log\left(\Pr\left(j|f,t\right)\right) = \\
&= \sum_{f,t}\sum_{j} n_{ft}^{j} \log\left(\sum_{d} \Pr\left(j|f,d,t\right)\Pr\left(d|f,t\right)\right) \\
&= \sum_{f,t}\sum_{j} n_{ft}^{j} \log\left(\sum_{d} \frac{\exp\left(\delta_{jft} + \beta_d \text{dist}(j,d)\right)}{1+\sum_l \exp\left(\delta_{lft} + \beta_d \text{dist}(j,d)\right)}\Pr\left(d|f,t\right)\right),
\end{aligned}$$

[24] We omit time variation, as these variables exhibit very little variation in our sample.

[25] Since our measure of length of berths or storage likely includes those for containers, we multiply berth length or storage space with the fraction of bulk ships that the port handles. The data on the percentage of bulk ships served is collected manually from MarineTraffic.

[26] We find that $\kappa_1 = 0.26$, which is 1.5 times the length of the average ship, and $\kappa_2 = 87.1$.

[27] Note that in our port call data we do not observe the domestic location, so we match observations at the month-foreign location level; we have a total of 560 such pairs and a choice among 51 ports.

where n^j_{ft} is the number of ships that choose port j in period t with foreign location f. We denote by $M_f = \sum_d M_{df}$ the number of potential shippers in market f, t, which we measure using the maximum observed yearly flow in FAF from/to foreign location f.[28]

Note that total time at port may be correlated with other time-varying characteristics of the port captured by ξ_{jft}. For example, an increase in demand driven by higher ξ_{jft}, by construction, leads to higher time at port, TT_{jt}, as more ships now choose port j. For this reason, it is necessary to use an instrument for TT_{jt}. We use unexpected disruptions in port j's service operations at t, which increase the total time a ship spends at port j. These disruptions directly affect TT_{jt} and, through the increased time at port, also reduce demand for port j as fewer shippers now choose this port; at the same time however, these disruptions affect how many ships arrive only through their impact on TT_{jt}, rather than directly. In particular, we employ the residuals from a regression of service time on a number of controls, including labor, cranes and fixed effects; these residuals are meant to capture unexpected disruptions in service, such as crane breakdowns and strikes.

In order to estimate the demand model using the above instrument, we follow Grieco et al. (2022). Their procedure, aimed at environments with micro-data on individual choices, but where endogeneity must be addressed, amounts to augmenting the likelihood above by a set of moments based on the exogeneity of ξ_{jft} from instruments, of the form

$$\mathbb{E}\left(\xi_{jft} z_{jft}\right) = \mathbb{E}\left[\left(\delta_{jft} - \beta_f \text{dist}(f,j) + \beta_T TT_{jt} + \beta_p p_{jt} - \gamma_t - \gamma_f - \gamma_{l(j)}\right) z_{jft}\right] = 0 \tag{10}$$

In addition, in order to identify the parameter β_d, which measures the importance of the distance from the unobserved domestic location d, we further augment the above micro-data likelihood by what we call a "macro-likelihood". This matches (yearly) aggregate flows between every foreign location f and every domestic location d, through port j:

$$\mathcal{L}^{\text{macro}}\left(\delta, \beta_d\right) = \sum_{\text{year}} \sum_{f,d,j} \left(\sigma_{fjd,\text{year}}\right)^{n_{fjd,\text{year}}}$$

where $n_{fjd,\text{year}}$ is the annual flow from f to d through j. It is important to note that this term is based on the annual flows reported in FAF, so it is a substantially aggregated version of the port call microdata.[29]

[28] It is worth noting that our measure of market size, M_f, is likely underestimated since it relies on the maximum number of shippers observed; it thus ignores potential shippers that have not yet engaged in imports/exports in this market.

[29] To be more precise, the FAF reports annual flows from every foreign location f to every domestic location d that go through a regional set of ports (e.g. ports in LA county). We use the micro-data to split total regional flow to its

Putting everything together, our estimation procedure minimizes the following objective function:[30]

$$arg \min_{\beta_\delta, \beta_d, \delta} \mathcal{L}^{\text{micro}}(\delta, \beta_d) + \frac{1}{2} m(\delta, \beta_\delta)' W m(\delta, \beta_\delta) + \alpha \mathcal{L}^{\text{macro}}(\delta, \beta_d)$$

where α is a weight for the "macro-likelihood", W is a weighting matrix, $\beta_\delta = \left(\beta_f, \beta_T, \beta_p, \gamma_t, \gamma_f, \gamma_{l(j)}\right)$, and $m(\delta, \beta_\delta)$ is the empirical analog of (10), so that,[31]

$$m(\delta, \beta_\delta) = \sum_{f,t} \sum_{j} z_{jft} \left(\delta_{jft} - \beta_f \text{dist}(f, j) + \beta_T TT_{jt-1} + \beta_p p_{jt} - \gamma_t - \gamma_f - \gamma_{l(j)}\right).$$

Finally, as described in Section 3.2, we enrich our demand model with a time-series process describing the evolution of the aggregate demand shocks γ_t. To do so, we leverage historical macro data. In particular, once we obtain the month fixed effects γ_t we project them on a set of macroeconomic variables y_t, including the Baltic Dry Index, commodity prices (specifically prices for coal, metals and agricultural commodities, collected from the IMF), and the Industrial Price Index:

$$\gamma_t = \beta y_t + \epsilon_t.$$

The benefit of introducing the observed macro series, y_t, is that we can now use a much longer time-series (1992-2023) based on these variables to estimate the properties of the macroeconomic conditions, $\hat{\gamma}_t \equiv \hat{\beta} y_t$. Once the parameters β is estimated, we obtain the parameters in (7) via Maximum Likelihood.

4.2.1 Results

Table 2 presents the estimated parameters. We begin by computing the "value of time" implied by our estimates, which is measured by the ratio β_T/β_p and captures the amount carriers are willing to pay in order to reduce time at port by one day. We find a value of time of \$0.9 per ton of deadweight, hence

corresponding port shares; if we were to use the flows through the broader region of ports we would lose concavity of the objective function and lose all the nice properties of the Grieco et al. (2022) approach.

[30] Standard errors are computed via bootstrap samples. In particular, we sample foreign region-month pairs with replacement. That the objective function is additively separable across markets implies the optimal δ_{jft} stays the same for a given β_δ . The resulting confidence intervals are based on 500 bootstrap samples.

[31] We refer the interested reader to Grieco et al. (2022) for details on the motivation and implementation of this approach. There are two differences from that approach here: (i) our context is simpler since we have the micro data on the population, while they allow the researcher to only have micro-data on a sample; (ii) we augment their likelihood by the "macro-likelihood". Importantly, it is easy to see that the "macro-likelihood" preserves the concavity property of the objective function that we optimize. We have experimented with different values of α and the results are robust.

Demand Estimates	
Average time at port	−1.44
	(0.49)
Port price per ton	−1.60
	(0.02)
Foreign distance (days of travel)	−0.71
	(0.15)
Domestic distance (log days of travel)	−1.2
	(0.05)
East Coast FE	-0.98
	(2.10)
Great Lakes FE	-15.89
	(4.57)
Gulf FE	-8.22
	(3.73)
Foreign location FE	Yes
Month FE	Yes
Instruments	Unexpected disruptions to port operations

Table 2: Demand Estimates. This table reports the estimated parameters of the payoff function (4). Standard errors are obtained from 500 bootstrap samples.

$\$45,000$ for the average ship.[32,33] This estimate for the value of time is substantial, especially compared to the direct cost of $\$14,000$ that the shipper needs to pay to hold the ship for an extra day. The difference between our estimate of value of time and this direct cost reflects the cost of port disruptions as they propagate through the supply chain.

In addition, we find that the elasticity with respect to time at port is large, as shown in Figure 6, suggesting that any disruptions affecting time at port are likely to have a large impact on demand. Moreover it is important to note that there is large heterogeneity in the elasticity across space which we return to below.

Furthermore our estimates suggest that, in addition to time at port, a port's geography and its fixed

[32] As a robustness, we have re-estimated demand on specific subsamples, e.g. based on ship size or commodity. When restricting the sample to the smaller ship sizes (Handysize and Handymax) for which we have enough observations to run the likelihood, we find that the value of time is somewhat lower, but in fact quite similar. Moreover, the value of time is somewhat lower for low value commodities, such as coal, and higher for high value commodities, such as grain and fertilizers. Allowing such characteristics to enter demand directly is left for future research.

[33] In practice we use the four-month rolling average of time at port. In robustness exercises, we find that increasing the horizon yields higher (longer-run) elasticities.

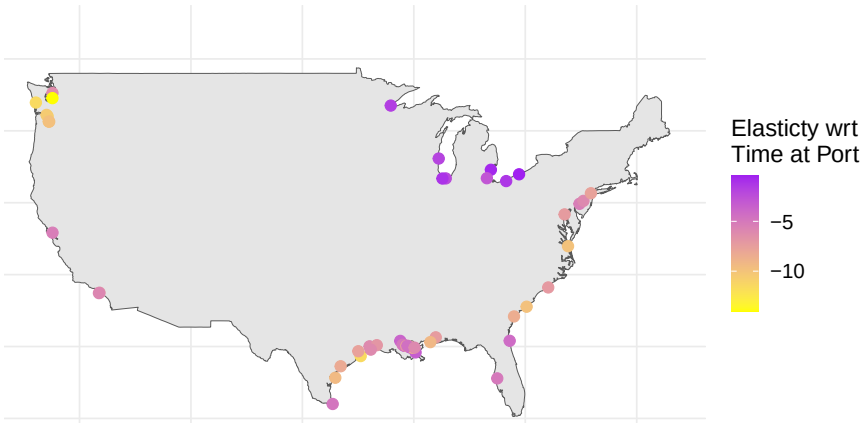

Figure 6: Demand Elasticity. This figure plots the average elasticity with respect to time at port for each US port.

characteristics also affect demand. In particular, our demand model allows for different forces that shape substitution patterns across ports, such as the distance to the foreign and domestic locations, as well as fixed port characteristics reflecting factors such as better services offered or better connectivity. For instance, ports in the East and West Coast have overall higher demand as captured by the region fixed effects in Table 2.

It is useful to understand the importance of geography vs. time at port. Figure 7 plots each port's geographic attractiveness, captured by the term $\delta^{\text{geo}}_{jfdt} = \beta_f \text{dist}(f,j) + \beta_d \text{dist}(j,d) + \gamma_{l(j)} + \xi_{jft}$, averaged across domestic and foreign locations, as well as time t. Note that $\delta^{\text{geo}}_{jfdt}$ depends purely on ports' geographic attributes, and purges variation coming from the ports' different congestion levels. The figure reveals that the geographically most attractive ports are ports in the Gulf, such as Houston (TX), that serve the flow of raw materials from the Mid Atlantic range and Southern states to Europe, Africa, and Latin America. Similarly, ports in the East Coast, such as Hampton Roads (VA) or Baltimore (MD), and Northwest Coast, such as Portland (OR), are also highly desirable, as they trade in raw materials with Europe and China.

Do traders use the geographically most convenient port? Interestingly, only 17% of market-shipper pairs choose the port that is the most desirable geographically, as defined by $\delta^{\text{geo}}_{jfdt}$ for that particular (d, f) pair. This suggests that time at port acts as a barrier to trade, pushing traders to sacrifice geographic convenience in order to avoid port congestion.

We further explore this by examining how shippers respond to an increase in time at port. Figure 8 plots the fraction of shipments that are diverted to a port's 10 closest neighboring ports, conditional on shippers choosing to switch. On average 37% of diverted trade is diverted locally. Moreover, again,

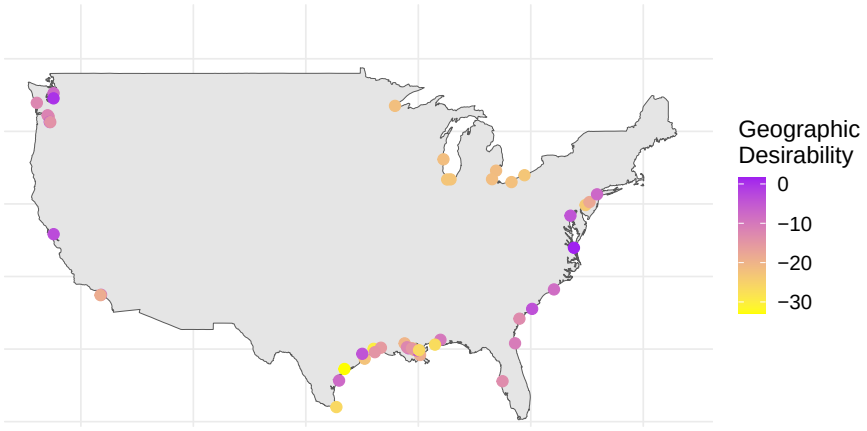

Figure 7: Geographic Attractiveness. This figure plots the variable $\delta^{\text{geo}}_{jfdt} = \beta_f \text{dist}(f,j) + \beta_d \text{dist}(j,d) + \gamma_{l(j)} + \xi_{jft}$, for each US port, averaged across domestic and foreign locations, as well as time t. Domestic and foreign locations are weighted by the market size of each (d, f) pair. This captures the geographic attractiveness of each port.

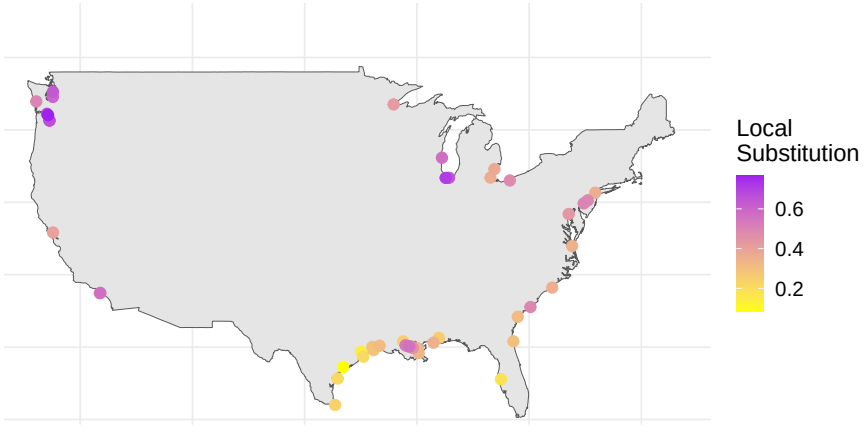

Figure 8: Local Substitution. For each port j this figure plots the fraction of shipments through j that is diverted away from j to a local port (i.e. to one of the 10 closest ports), conditional on shippers choosing to switch following a small increase in time at port.

there is large heterogeneity across ports. In general, the figure shows that for ports in the Great Lakes substitution tends to be local, while the reverse is true for ports in the East Coast and the Gulf.

It is instructive to compare two ports displaying these opposite patterns, Corpus Christi (TX), where 21% of trade is diverted locally, and Cleveland (OH), where 56% of trade is diverted locally. Corpus Christi (TX) is very attractive for a number of (d, f) pairs that are often far away, mostly serving trade between Europe and Latin America, and the Atlantic Coast and Midwest. Therefore, when Corpus Christi (TX) is congested, these importer-exporter pairs respond by substituting to Hampton Roads (VA), a large port on the East Coast, well connected with these origins and destinations. In contrast, Cleveland (OH) mostly serves local shippers, i.e. it is attractive mostly to local (d, f) pairs, operating between Canada and the Midwest. Naturally, these trades have less flexibility and tend to rely on other local ports.

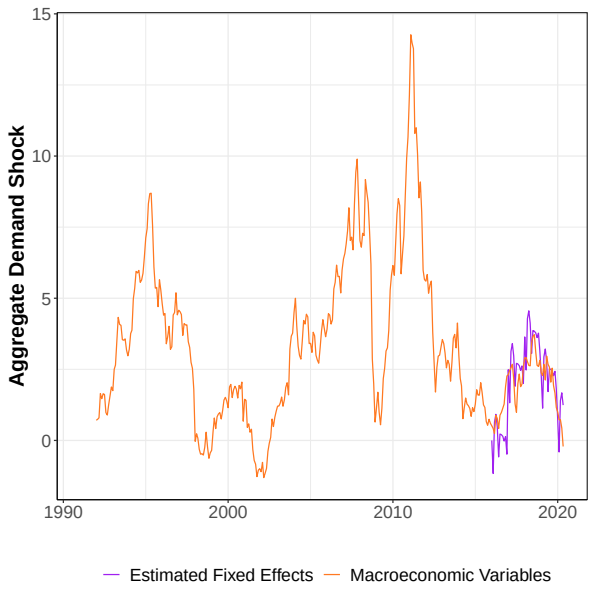

Figure 9: Macroeconomic volatility. This figure plots the estimated fixed effects, γ_t, as well as the predicted time fixed effects by macroeconomic variables, $\hat{\gamma}_t \equiv \hat{\beta} y_t$. The macroeconomic variables y_t include the Baltic Dry Index, commodity prices, and the industrial price index. See Table 9 in Appendix A for details.

These substitution patterns are helpful in understanding the large geographic heterogeneity in elasticities with respect to time at port, shown in Figure 6. For instance, shippers in the Great Lakes are fairly far away from the most desirable ports and when they substitute they tend to do so locally. Having said that, precisely because there are no good options nearby, the estimated elasticities of the Great Lakes are the lowest in the country. Conversely, ports in the Gulf or the East Coast have several substitutes in a wider geographical range, so they do not substitute locally and they feature high elasticities. Finally, ports in the Northwest Coast are close to some of the most attractive ports in the country, so they have both high elasticities and a much higher propensity to substitute locally when they do substitute.

Finally, Figure 9 plots the estimated time fixed effects, γ_t, which capture macroeconomic volatility, while Table 9 in Appendix A reports the estimated coefficients for the AR(1) process for $\hat{\gamma}_t \equiv \hat{\beta} y_t$ given in (7), as well as the estimated parameters β that allow us to extrapolate to a longer time-series. Figure 9 demonstrates a few interesting empirical patterns: (i) the fit during our sample is quite good with respect to the overall timeseries; (ii) our sample period is one of overall low demand for port services; (iii) macroeconomic volatility is quite severe.

5 Welfare Gains from Infrastructure Investment

Billions are spent every year from both public and private actors for investment in port infrastructure. What are the welfare gains of these investments? We now use our estimates for port technology and demand for port services to answer this question. In the next section, we compare these welfare gains to the infrastructure costs to obtain a measure of net returns.

We compute the aggregate welfare gains to an increase in the capacity of the US port system by one slot. In other words, a single port can now handle one additional ship; this is similar to the approach followed in Allen and Arkolakis (2022). To put this in context, this increase corresponds to approximately a 1% increase of the total US capacity. In practice, we increase K_j to $K_j + 1$ for a given port j and then we repeat this exercise for all J ports, one port at time. We use our model to simulate a long-run counterfactual time series of 30 years. We assume that along with K_j, port j's labor and cranes also increase so as to keep its (average) service time constant.[34] Note that since many small ports can handle only 1 or 2 ships at a time this increase in capacity can be substantial. We describe how these simulations are conducted in Appendix C.

Mechanisms We begin by outlining the economic mechanisms that are set in motion when the capacity of a port increases. We first describe the impact at the "treated" port, and then explore what happens to all other ports.

When capacity K_j increases at a **treated port** j, its wait time declines, all else equal. In turn, the reduction in congestion makes the port more appealing and attracts ships both from competing ports and from the outside good. It is important to note that the latter creates a net increase in total trade, unlike the ships' reallocation, which is a transfer of trade from other ports to port j.

Furthermore, the endogenous response of demand at port j dampens the original fall in congestion. Indeed, although congestion initially declines because of the higher capacity, the resulting higher demand mitigates this decline, boosting congestion and time spent waiting. Therefore, while increasing K_j to $K_j + 1$ reduces port j's wait time, this reduction is mitigated by the increased arrival of ships that respond to this change. This is reminiscent of the so called "Iron Law of Congestion" (Downs, 1962, 1992; Duranton and Turner, 2011) often evoked in road construction: as road capacity increases, new drivers

[34] This is done in order to recognize that to some extent additional berth space may not increase effective capacity without adding resources to service this additional vessel. A different exercise would be to keep labor and cranes fixed and only change K_j. While less costly, it would also lead to higher service times, as the same inputs would now need to be spread over more berths and this increase in service times will be larger at smaller ports that have fewer inputs. We consider our approach a more realistic scenario.

use the road, mitigating any relief in congestion. In our context, this inflow of new customers is socially desirable as it potentially constitutes new (international) trade that was not taking place before because of port congestion.

In summary, we expect to see a potentially considerable increase in trade at the treated port (depending on the port's attributes, an issue we return to below), but a more modest decline in congestion. Indeed, we find that on average (median) trade increases by 42% (38%) at the treated port, while congestion declines on average (median) by 4.1% (3.7%).

Now let us turn to what happens to **all other ports**, when capacity increases at port j. In this interconnected system, changing capacity in one port does not occur in a vacuum. Instead, it is bound to have an impact on the entire port system, as suggested by the above discussion. Improving port j does not affect just j but also other ports, whose demand is reduced as their competition now offers better wait times. In other words, all other ports lose demand because of substitution to the treated port. This *substitution effect* implies that the increase in trade at the treated port does not necessarily increase aggregate trade.

To illustrate this feedback effect with an example, we consider an increase in the capacity of Hampton Roads (VA), holding time at port fixed at other ports for now. As shown in the left panel of Figure 10, this change can substantially reduce ship arrivals at competing ports; indeed several ports in the Gulf and East Coast see their demand decline up to 7%, all else equal.

But the story does not end here. As all these other ports lose demand, they become de-congested, which in turn endogenously boosts their demand. This feedback effect in neighboring ports creates positive spillovers: investing in infrastructure in port j can decongest not just j, but also a wider set of ports, especially those that are close substitutes of j. We illustrate this in the right panel of Figure 10, which plots the full effect of increasing capacity at Hampton Roads (VA). Interestingly, this feedback effect can be quite large: for several ports, such as Baltimore (MD) or Charleston (SC), the full effect is almost null, despite a considerable substitution effect. In other words, although a large share of exporters substituted away from Baltimore (MD) towards Hampton Roads (VA), time at port at Baltimore (MD) fell sufficiently to re-attract demand to the point that overall Baltimore (MD) did not lose much traffic. In this case, the overall effect on Baltimore's trade is close to zero, as the two forces roughly offset each other.

In summary, at all other ports, we expect to see some decline in trade, as well as some decline in congestion. Indeed, we find that at all other ports on average trade decreases by 0.19%, while congestion declines on average by 0.6%.

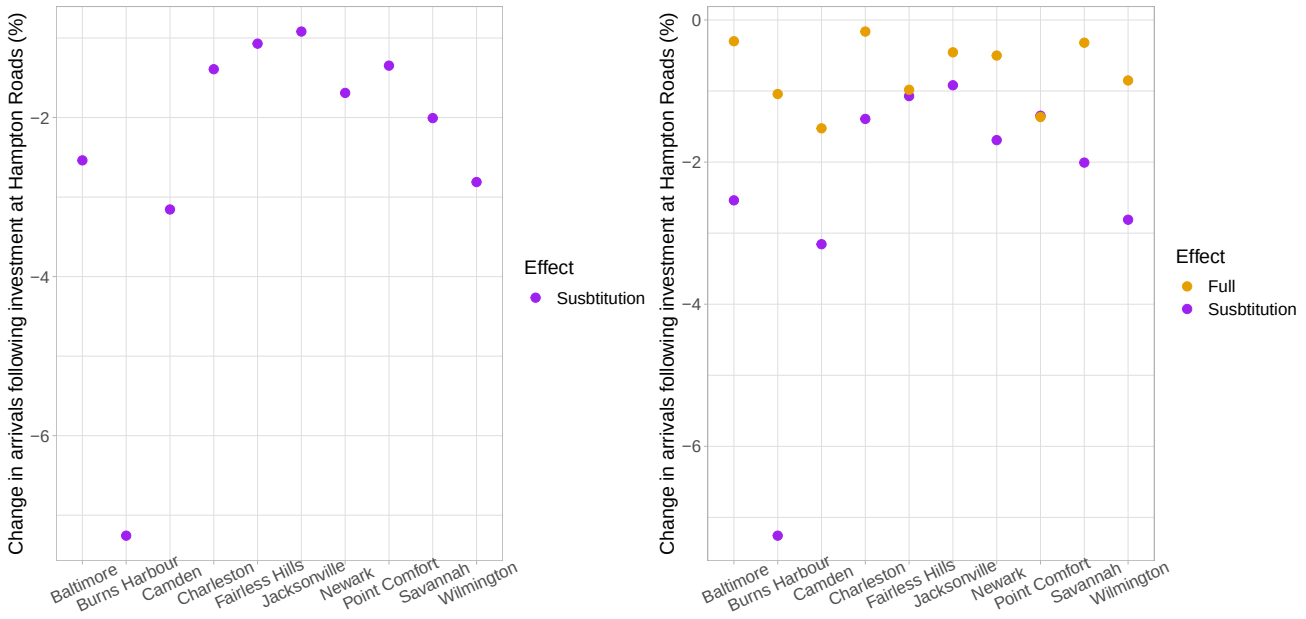

Figure 10: Substitution and full effect of increasing port capacity at Hampton Roads (VA) from K_j to $K_j + 1$, so that it can now handle one additional ship. The left panel plots the change in arrivals at ports other than Hampton Roads (VA), holding time at port at its original level there. The figure depicts the port's top 10 substitutes. The right panel plots the full effect on other ports' demand in response to an increase in infrastructure in Hampton Roads (VA) (i.e. not holding time at port constant at other ports).

The previous discussion has focused on changes in trade flows to different ports. More generally, changes in aggregate welfare reflect increases in aggregate trade, as well as gains stemming from shipper reallocation. For instance, shippers who now choose to use the port of Hampton Roads (VA) instead of Baltimore (MD) might benefit from closer proximity to their final destination, or the lower congestion. We next discuss in detail how investment in infrastructure affects aggregate welfare.

Trade and Welfare Gains The key outcome a policymaker is interested is not congestion per se, but rather the gains in trade and welfare that are generated by new infrastructure. As described above, in order to evaluate the welfare impact of capacity improvement in a single port, we need to also take into account the overall effects across all ports. Figure 11 depicts this total change in trade for *the entire* US port system when we increase the capacity of one port at a time; Figure 12 depicts the corresponding gross welfare gains to infrastructure investment, which in our setup is the change in *total* welfare from changing a port's capacity. Table 10 in Appendix A contains the change in wait time, trade and welfare from expanding capacity K_j to $K_j + 1$ for every port in our data.

We find that, on average, investing in infrastructure increases aggregate trade by 0.67% and aggregate welfare by 0.5%. Importantly, however, the returns are quite heterogeneous across ports. Indeed, trade

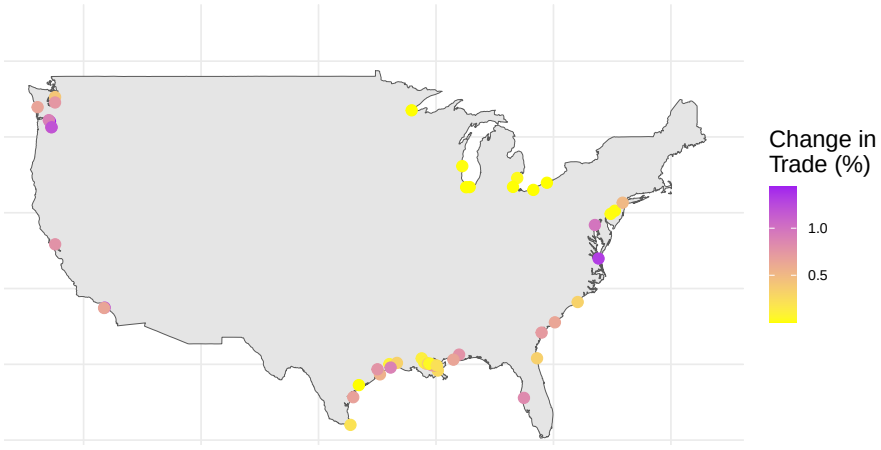

Figure 11: Trade impact of investment in infrastructure. In this figure, we increase port capacity of port j from K_j to $K_j + 1$, for all j one at a time, so that port j can now handle one additional ship. The figure depicts the total change in trade from increasing capacity in each port j, one at a time.

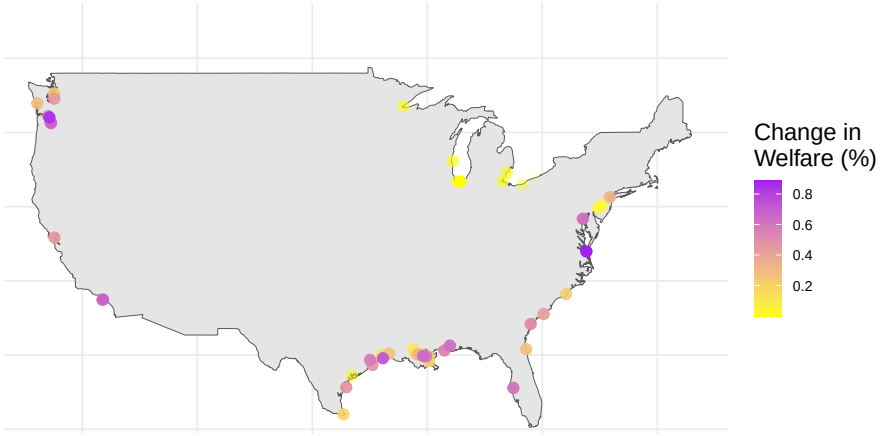

Figure 12: Welfare impact of investment in infrastructure. In this figure, we increase port capacity of port j from K_j to $K_j + 1$, for all j one at a time, so that port j can now handle one additional ship. The figure plots the total change in welfare from increasing capacity in each port j, one at a time.

and welfare can increase by more than 1% when investing in ports such as Hampton Roads (VA); but they are negligible when investing in other ports, such as the ones in the Great Lakes.[35] This suggests that targeting investment to the "right" port, as opposed to indiscriminate investment, can increase the returns to investment manyfold.

In addition, our estimates suggest that spillovers are large. In particular, if we only allow for the substitution effect, the aggregate increase in trade is 20% lower on average, and up to 39% lower for the ports with the highest gains.[36]

What makes a port generate high (vs. low) welfare gains? First, size alone does not directly predict

[35] These numbers are in line with those found in the transport infrastructure literature, such as Fajgelbaum and Schaal (2020) and Allen and Arkolakis (2022).

[36] As above, to isolate the substitution effect we compute the change in trade holding time at port at other ports fixed.

welfare gains from infrastructure investment. As an example, Hampton Roads (VA) is the sixth largest port in the US, but delivers the highest gains. To explore the drivers of welfare gains from investing in infrastructure, we consider how these gains depend on both the characteristics of the treated port, as well as those of the set of ports it decongests. Consider again Hampton Roads (VA). Following our discussion in Section 3.1.1, an increase in K_j should have the largest impact on ports that are already congested: investing in infrastructure in ports with spare capacity has little effect; on the contrary, an additional slot in ports that have low capacity relative to their demand, can have a substantial impact, as higher K_j helps absorb the high demand. As shown in the left panel of Figure 13, Hampton Roads (VA) is indeed a congested port.

In addition, the geographic attractiveness of the treated port is bound to be decisive for its welfare gains. Indeed, investing in ports that are in a convenient location can both attract more shippers from the outside option, and benefit existing shippers that reallocate there. As shown in the right panel of Figure 13, Hampton Roads (VA) is the most attractive port in terms of the variable δ^{geo} , which primarily captures the proximity of the port to diverse origins and destinations.

However, looking at the treated port alone does not suffice to break down the overall welfare gains it generates. We must also consider the set of ports that obtain the highest spillovers, i.e. the set of ports that the treated port decongests. In line with our discussion above, if these are ports that can themselves generate large gains, the overall impact will be higher. Figure 13, demonstrates that the 10 closest substitutes of Hampton Roads (VA), are themselves both fairly congested, and geographically attractive.[37]

This intuition holds more generally. As shown in Table 11 in Appendix A, welfare gains are higher for ports that have higher average levels of congestion, captured by the average time that ships have to wait in anchorage, and for ports that have higher geographic attractiveness, δ^{geo} . Moreover, gains increase in the average congestion and δ^{geo} of the treated port's closest substitutes. As discussed in Section 4.2, congestion acts as a barrier to trade; investing in ports that are geographically convenient ensures they are accessible to a larger number of shippers, lowering the transportation costs that they need to shoulder.

The presence of the spillovers we document raises important policy questions regarding the coordination of investment activities. An implication of our analysis is that unless decisions are made centrally, they are unlikely to internalize the spillovers across ports. While some port investment decisions are made at the federal level, most of the funding comes at the private or local (e.g. city, state) level. For

[37] We identify these ports as those with the highest cross derivative with respect to time at port at Hampton Roads (VA).

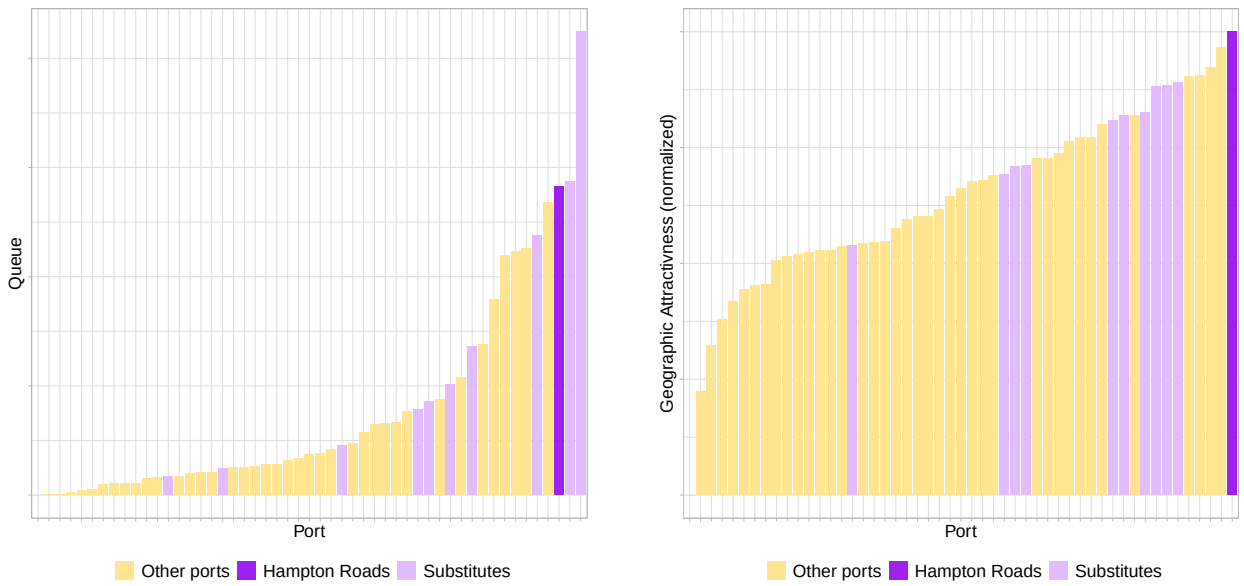

Figure 13: What makes a port generate high welfare gains? The left graph ranks ports by their queue. Hampton Roads (VA), the port with the highest returns, is fairly congested. In addition, its 10 closest substitutes are themselves also fairly congested. The right graph ranks ports by their geographic attractiveness, δ^{geo} . Hampton Roads is the most geographically attractive port and its 10 closest substitutes are themselves highly attractive ports.

instance, the Biden administration's 2021 infrastructure bill allocates over $2.25 billion for ports over the next decade;[38] however this is only a small fraction of the $160 billion to be spent between 2021 and 2025. Given that local authorities seem to care predominantly about the port's local effects, one may be concerned that these spillovers are not fully internalized. Indeed, this lack of coordination and competing incentives might explain why, as we will see in the next section, for some ports the net returns from investment can be sizable.

6 Net Returns of Investment in Port Infrastructure and Volatility

When and where is port investment justified? How should a budget be allocated? How frequent do macroeconomic shocks need to be to warrant investment? To answer these questions we calculate the net returns from investment in port infrastructure. While the welfare gains were computed in Section 5, the cost of port infrastructure is difficult to obtain; as discussed in Section 2.2, we rely on cost estimates from the US Army Corps of Engineers (USACE), as well as land costs required for storage. It is important to note that authorities rely on the same information from the USACE to make investment decisions.

We first describe how we calculate the cost of port infrastructure (Section 6.1); then we present the

[38]See Section Port Infrastructure Development Program on page 1442 of the "Infrastructure Investment and Jobs Act" (Public Law 117–58, November 15, 2021).

net returns and we explore the role of macroeconomic volatility (Section 6.2).

6.1 The Cost of Expanding Port Capacity

We calculate the cost of increasing port j's capacity from K_j to $K_j + 1$, so that the port can handle one additional vessel. The cost for port j is equal to,

$$C_j^d + C_j^l + C_j^b$$

where C_j^d is the cost of dredging a section of ocean or riverbed at port j sufficient to accommodate the ship; C_j^l is the cost of purchasing a parcel of land bordering port j to serve as a storage facility for goods transiting the new berth; and C_j^b is the cost of constructing a bulkhead to support the new berth.[39] We discuss each term in turn. Appendix D provides the details.

Accommodating an additional vessel requires dredging a portion of the seafloor with volume determined by the vessel's size and the port's current depth; the associated cost, C_j^d, crucially depends on what share of this volume requires blasting through rock. For each berth in our dataset, the USGS provided us meter-by-meter estimates of underwater wave velocities, computed based on Boyd (2020).[40] From these, we calculate the depth at which seabed switches from rock material to sediment, whereafter drilling and blasting must be used. In other words, we obtain an estimated depth to non-dredgeable rock material for each berth.

We estimate the cost function of dredging rock and non-rock material from cost estimates obtained by the USACE. In particular, we are able to use data from two expansion projects: Delaware in 2010 and Boston in 2017. For these projects, the USACE provides detailed estimates of predicted dredging costs and material quantity for 78 berth-depth combinations. We estimate a linear function of costs on total quantity dredged (for both rock and non-rock; the total cost is the sum of the two).

Next we turn to the cost of storage, C_j^l. This is equal to the price of industrial land around port j, p_j^A, times the acreage A_j required to accommodate the additional ship, i.e., $C_j^l = p_j^A A_j$. To obtain land prices around ports, we first use OpenStreetMap to isolate industrial and commercial land in proximity to each berth. Then, we find the corresponding value of this land from data made public by Nolte (2020). We obtain the acreage required to accommodate the additional ship from the Leontief capacity function

[39] A bulkhead is a vertical shoreline stabilization structure that primarily retains soil, and provides minimal protection from waves.

[40] We are most thankful to Oliver Boyd for sharing this data with us.

(9).

We also need a measure for the cost of bulkhead construction, C_j^b. The USACE reports that the cost of erecting a new bulkhead at Charleston (SC) amounted to $22 million, or $6,600 per foot of berth. We use this number as our baseline and perform robustness around it.

Finally, we assume that, when investing in port infrastructure, the port increases labor and cranes in order to maintain the service time constant and we add the cost of purchasing cranes and hiring workers to the total expansion cost. To recover the number of cranes and workers required, we use the CES service time specification (8). We calibrate the cost of a crane to 20 million USD, based on estimates from the port of Savannah (GA), and the annual wage at $\$60,000$ (US Department of Labor).

Overall, total costs are 115 million USD on average with a fair amount of geographical dispersion (standard deviation equals 98 million USD). Land costs account for about 49% of total costs on average, and dredging costs for 26%. Land costs account for an even larger share of total costs in urban areas such as Portland (OR) (87%), Seattle (WA) (85%), and New Orleans (LA) (80%).

Finally, let us note that there are two caveats to this data. First, they rely on two port projects (albeit with several berths each), as these are the only projects reporting detailed information on costs and quantities excavated.[41] Second, they rely on engineering estimates, while economists often worry about other types of costs that are relevant for investment decisions (such as opportunity costs, disruptions to port operations, political constraints etc.). Importantly, however, policy makers rely on the same cost data, which they benchmark against a (local) benefits analysis of increased port throughput; one can think of our exercise as comparing the same costs to a model for "benefits" that takes into account dynamics, demand feedbacks, and spillovers across ports.[42]

6.2 Net Returns and Volatility

In Figure 14 we compare the welfare gains from increasing K_j to $K_j + 1$, calculated in Section 5, to the corresponding costs, and highlight in purple (yellow) the ports for which the former are higher (lower) than the latter. Net returns are positive for 15 out of 51 ports, and are highest for the ports in the Gulf, such as Convent (LA), and Corpus Christi (TX), as well as for Hampton Roads (VA). Both welfare and costs shape up the net returns: as shown in Figure 20 in Appendix A, although the ports with the

[41] Reassuringly, the estimated cost function across the two projects is very similar; if we only use each project in isolation to estimate the dredging cost function, we obtain very similar results.

[42] An alternative approach would be to estimate the cost of infrastructure from port investment decisions, following the IO literature. There are two important issues with this approach that made us uncomfortable using it. First, this requires a model of port behavior; yet it is quite difficult to gauge what the objectives of port authorities are. Second, investments in port expansion are infrequent in the US.

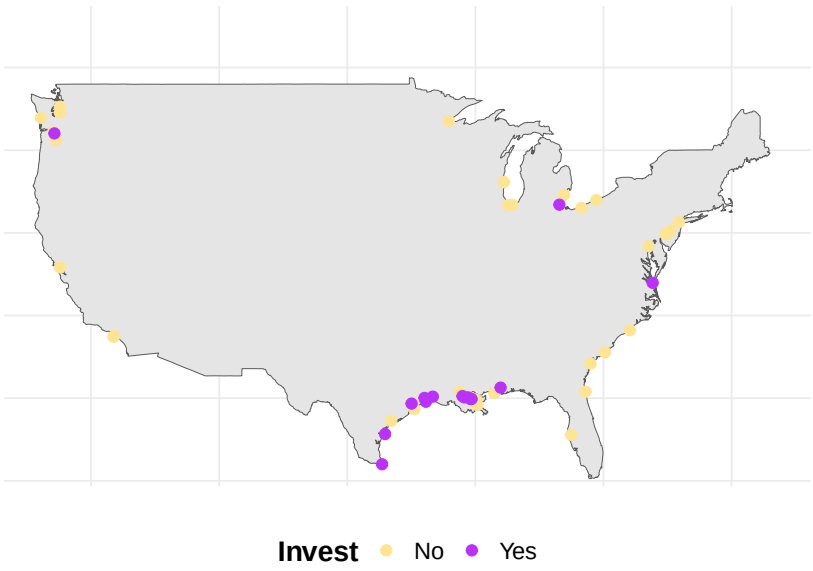

Figure 14: Net returns of investment in infrastructure. In this figure, we increase port capacity of port j from K_j to $K_j + 1$, for all ports, one at a time, so that port j can now handle one additional ship. The graph highlights in purple the ports where the returns from investment are higher than the costs, and yellow for the reverse.

highest welfare gains also generate positive net returns, there are several ports, which may be yielding high welfare gains, but also exhibit prohibitive costs, turning net returns there negative. For instance, the ports of Long Beach (CA) and Baltimore (MD) have some of the highest welfare gains from investment, but are both located in urban areas which are associated with a steep cost of land expansion. The last column in Table 10 in Appendix A contains the rate of return from expanding capacity K_j to $K_j + 1$ for every port in our data.

An important innovation of our framework with respect to the literature is that it explicitly considers congestion dynamics, which illustrates that infrastructure investment is much more valuable when demand surges. We can thus address a timely policy question: how frequent do macroeconomic shocks need to be to warrant investment? In recent years, there is increasing evidence that shocks are becoming more frequent and that uncertainty is rising (see Bloom, 2009, as well as the literature on the recent supply chain disruptions); at the same time policymakers appear increasingly concerned about resilience.[43]

In Figure 15 we explore net returns from investment in infrastructure as the frequency of macroeconomic shocks gradually rises.[44] Following our discussion in Section 3.1.1, for the average port, welfare gains increase steeply with volatility, and they double as the standard deviation of shocks doubles, as shown in the left panel of Figure 15. Even more strikingly, at the 90^{th} percentile, welfare gains quadruple.

[43] Indeed, the World Uncertainty Index has roughly doubled in the last twenty years.

[44] In practice, we repeat the long-run simulations under a mean-preserving spread of the aggregate demand process (i.e. we increase the standard deviation, σ_ϵ, of the aggregate demand γ_t). Of course, our estimated net returns would be amplified if the mean of the demand process also were to increase.

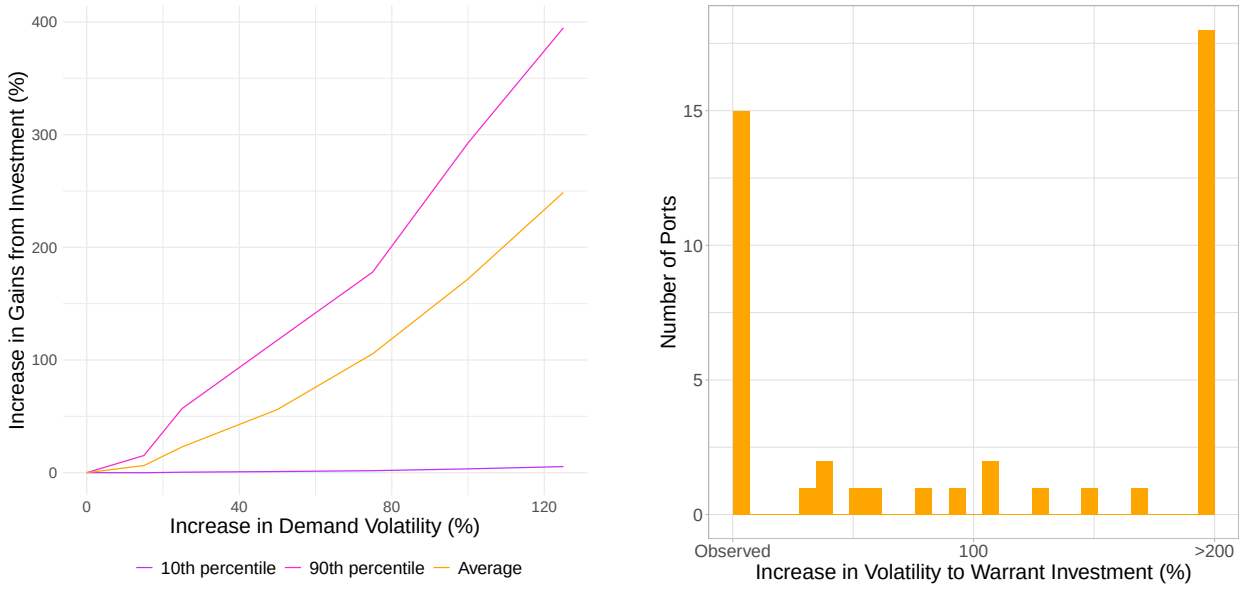

Figure 15: Rising macroeconomic volatility. The left panel plots the welfare gains from investment in the average, as well as the 10^{th} and 90^{th} percentile of US ports, for different percent increases of aggregate volatility. In particular, we repeat the long-run simulations under a mean-preserving spread of the aggregate demand process (i.e. we increase the standard deviation, σ_ϵ, of the aggregate demand γ_t). The right panel plots the histogram of the increase in aggregate volatility that turns net returns positive for different ports.

At the other extreme, at the 10^{th} percentile, welfare gains are barely affected by the volatility increase. In fact, this remains true until the 25^{th} percentile.

The right panel of Figure 15 plots the increase in aggregate volatility that makes net returns positive for different ports. As already discussed above, for 15 ports net returns are already positive at the observed volatility. At the other extreme, for a sizable share of ports (25%) net returns do not turn positive unless volatility more than triples; for these ports, investment is unlikely to be warranted even if the policymaker only cares about resilience. Finally, there is an intermediate range of ports for which some increase in volatility makes investment valuable; for instance, 10 additional ports yield positive net returns if volatility doubles.

Interestingly, increasing macroeconomic volatility changes not just the level of returns, but also their geography. Indeed, the ports that become critical as volatility increases are not necessarily those with the highest baseline returns. In particular, as shown in Figure 16, many of the additional ports whose net returns turn positive when volatility doubles are concentrated in the Great Lakes, where welfare gains increase six-fold on average. Ports in that region are close substitutes and their aggregate capacity can normally absorb occasional demand shocks without wait times exploding. However, when extreme demand shocks hit, all the ports there become overwhelmed. Due to the limited ability for substitution outside the region, increasing capacity is the only way to alleviate congestion in the area.

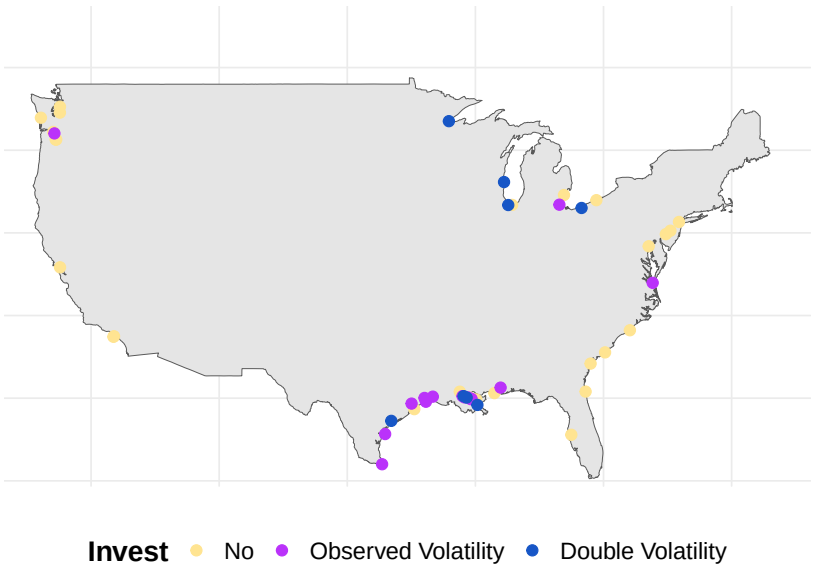

Figure 16: Rising macroeconomic volatility. This figure depicts the map of net returns: purple ports are the ports with positive net returns at the baseline level of volatility; blue ports are the ports whose returns turn positive when volatility doubles; and yellow ports are ports whose net returns remain negative.

Investment in Multiple Ports The analysis so far considers changes in one port at a time, but in principle policymakers might be interested in investing in multiple ports. We repeat our analysis assuming that we can invest in multiple ports simultaneously and examine how the results change. In particular, as an initial step towards this complex problem, we consider the net returns of expanding capacity from K to $K+1$ at 3 ports simultaneously, for the 10 ports with the highest net returns (i.e. we compute net returns for 120 port combinations). We find that net returns remain positive for all these combinations, suggesting net returns at the top ports are not marginal. More interestingly, when exploring the five combinations that yield the highest returns, we note that all but one involve one port in the Gulf, one in the East Coast and one in the West Coast, suggesting that it might be optimal to "spread out" investment throughout the country. This is consistent with the presence of spillovers, that the policymaker can harness: investing in one port decongests an area, making additional investment far away more valuable.

7 Conclusion

In recent years, infrastructure investment has been in the pubic spotlight. Ports in particular have undergone a series of major disruptions, from Covid-19, to chokepoint closures. How costly are these disruptions? In this volatile environment, what are the returns to investment in infrastructure? How should a budget be allocated? How frequent do macroeconomic shocks need to be to warrant investment? Answering these questions requires a clear understanding of congestion formation, shippers' response to

shocks, and the interdependencies inherent in the port network. This paper develops simple tools to address these issues. We showcase that disruptions are significantly costly. Yet, for investment to be justified, it must be coordinated and targeted. This is because of the substantial heterogeneities across ports, as well as the sizable spillovers generated by investment. Finally, as these are long-run investments, their returns are crucially affected by macroeconomic volatility; the increased frequency of shocks, or the stronger interest in resilience, will make policymakers more willing to invest.

Appendix

A Additional Tables and Figures

Dependent Variable:	Dues per Port Call (log)			
	(1)	(2)	(3)	(4)
Deadweight (log)	0.6169***	0.6176***	0.6143***	0.6218***
	(0.0445)	(0.0435)	(0.0437)	(0.0255)
port FE	Yes	Yes	Yes	
loading/unloading FE	Yes	Yes	Yes	Yes
cargo FE	Yes	Yes	Yes	Yes
year FE		Yes		
month FE			Yes	
port-year FE				Yes
N	24,119	24,119	24,119	24,119
R^2	0.687	0.688	0.690	0.693

Table 3: Variation in port charges. This table presents a regression of port prices on ship characteristics, as well as time and port fixed effects.

	log (num)	log (hrs)
log (members)	1.031***	0.980***
	(0.097)	(0.104)
N	81	81
R^2	0.589	0.530

Table 4: Correlation between union membership and annual man-hours for ports on the West Coast using shift-level data from the PMA (Pacific Maritime Association). We regress the number of total workers (log(num)) in shifts in the detailed data from the PMA, as well as total hours (log(hrs)) on the total number of union members. The estimates suggest that not only is the correlation high, but the elasticity is equal to one.

	Service time (log)	
# of other ships serviced (log)	0.72^{***}	0.73^{***}
	(0.076)	(0.072)
queue (log)		-0.025
		(0.020)
N	59,576	59,576
Port FE	Yes	Yes
Commodity FE	Yes	Yes
Month-year FE	Yes	Yes
Ship size	Yes	Yes

Table 5: Service time and number of ships at port. In this table we regress port service time on the number of ships simultaneously handled, as well as the number of ships in anchorage, controlling for port, commodity, and month-year fixed effects. Service time per ship depends on how many other ships are serviced at the same time, but not on the queue of ships waiting.

	Total workers per shift
n_{bulk}	31.21^{***}
	(2.49)
$n_{container}$	83.67^{***}
	(1.88)
$n_{container} \times n_{bulk}$	0.14
	(0.42)
N	27,251
R^2	0.86
Port FE	Yes
Year-Month FE	Yes

Table 6: Interaction between bulk ships and containerships. We use shift level data for West Coast ports from the PMA (Pacific Maritime Association) to test for interactions between bulk and containership operations within a port. We find that the total number of workers per shift increases as a higher number of containerships ($n_{container}$) or bulk ships (n_{bulk}) are serviced. However, the interaction is not economically and statistically significant, suggesting that bulk and container loading and unloading operations are largely independent.

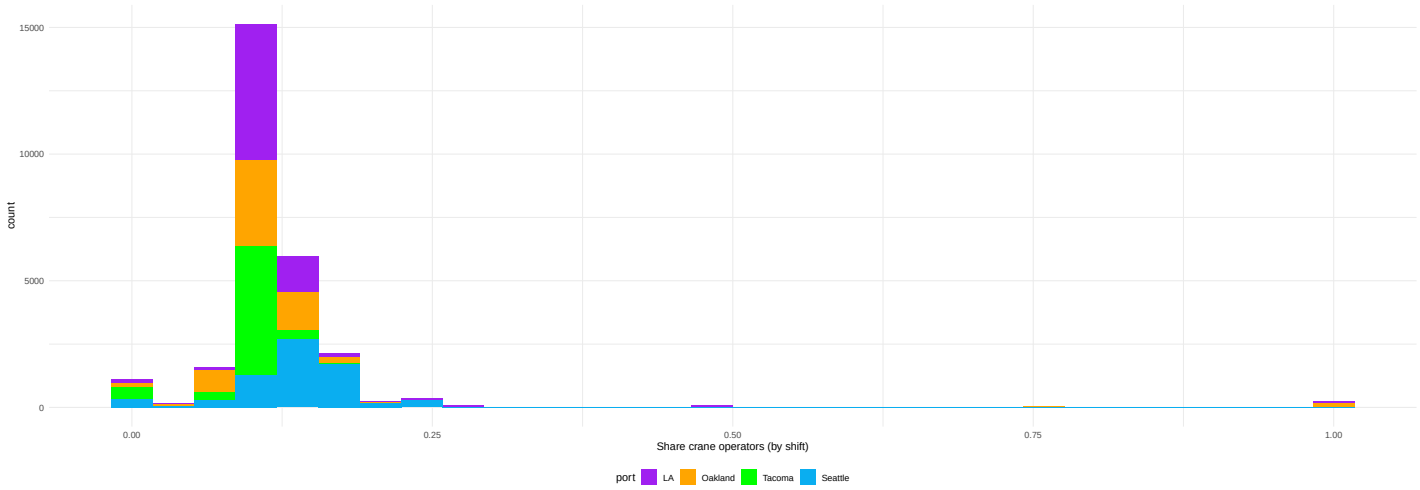

Figure 17: Is crane technology substantially different across ports? This figure uses granular labor data for the major West Coast ports from the PMA (Pacific Maritime Association) and plots the histogram of the share of crane operators with respect to all port workers and illustrates that it does not vary much across ports and over time.

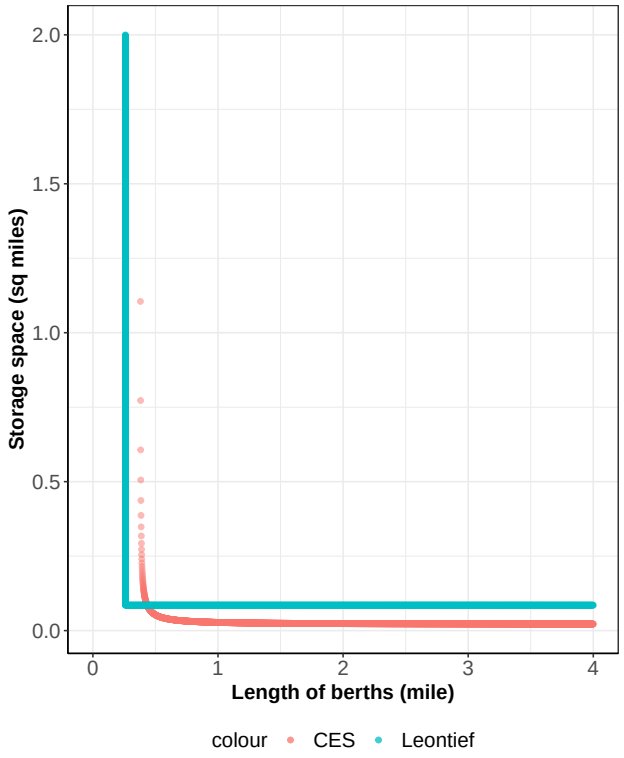

Figure 18: Capacity function. This figure compares the calibrated Leontief capacity function against an estimated CES capacity function and shows that the two are similar. In particular, we estimate $K_j = z_j \left[\alpha \left(\frac{B_j}{\kappa_1}\right)^{-\gamma} + (1-\alpha)\left(\frac{A_j}{\kappa_2}\right)^{-\gamma}\right]^{-\frac{1}{\gamma}}$ via GMM. We rewrite it as $\log K_j = \log \bar{z} + \epsilon_j - \frac{1}{\gamma}\left(\alpha \left(\frac{B_j}{\kappa_1}\right)^{-\gamma} + (1-\alpha)\left(\frac{A_j}{\kappa_2}\right)^{-\gamma}\right)$ and use the following two moment conditions: $E\left[B_j \epsilon_j\right] = 0$ and $E\left[A_j \epsilon_j\right] = 0$ to obtain the parameters.

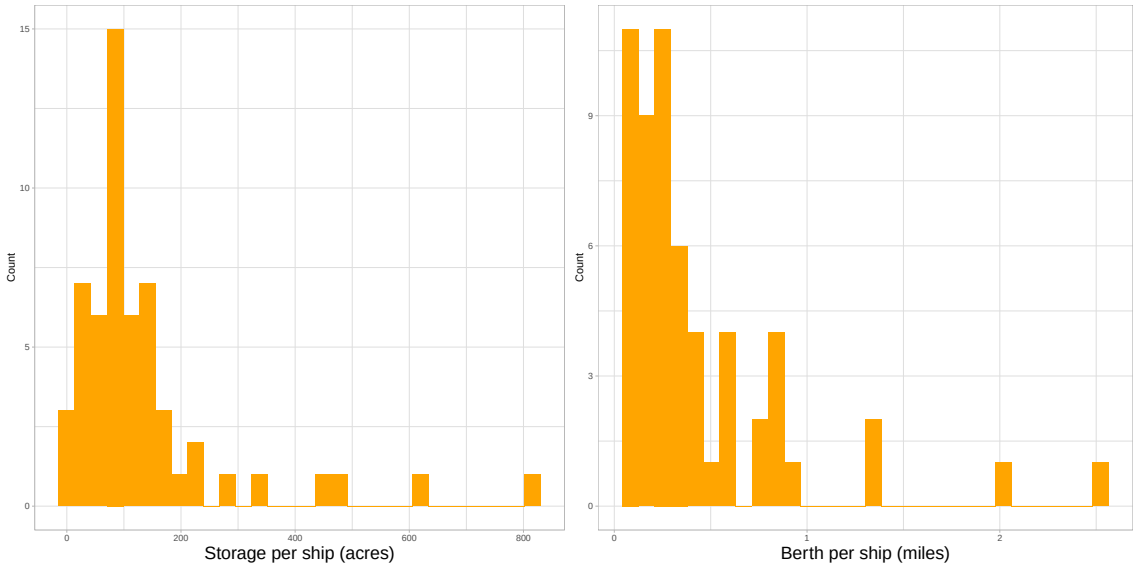

Figure 19: Storage space and berth length per ship. This figure plots the histogram of the ratio of storage space per ship across ports, i.e. A_j/K_j, as well as the ratio of length of berth per ship across ports, i.e. B_j/K_j.

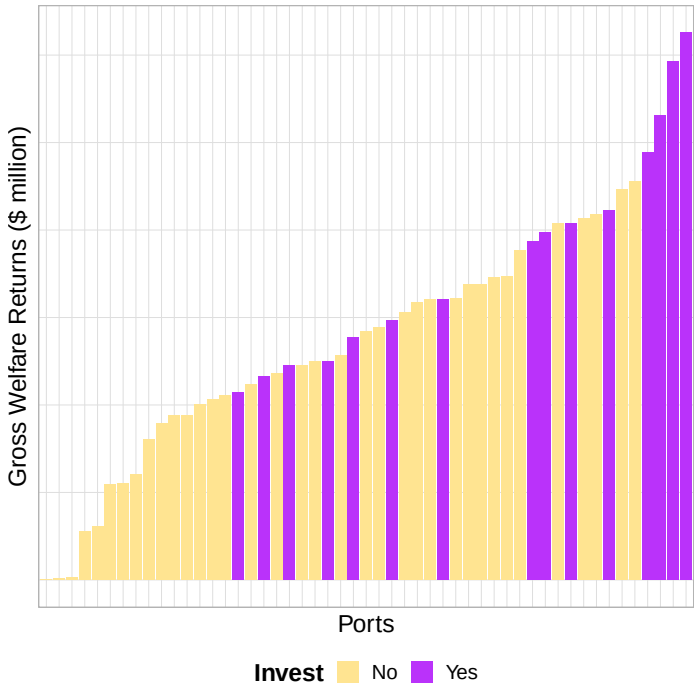

Figure 20: Welfare gains from investment in US ports. Ports are ranked according to their calculated welfare gains. The color varies depending on whether the gains are higher (purple) or lower (yellow) than the cost.

	ω_{jt}
Tidal range	0.231
	(0.076)
Port Type: Coastal natural	−0.255
	(0.075)
Port Type: River	−0.02
	(−0.06)
Port Type: Coastal breakwater	0

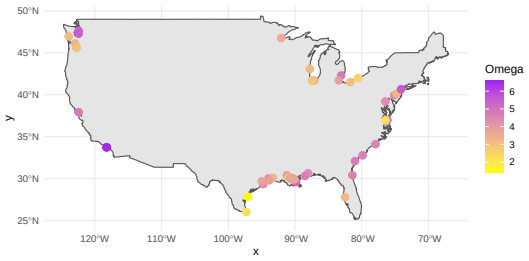

Table 7: Port Productivities. The table on the left presents the results from a regression of the estimated port productivity on tidal range and an indicator for port type obtained from the World Port Index (WPI). Coastal natural ports are sheltered from the wind by virtue of their location within a natural indentation; coastal breakwater ports lie behind man-made breakwater constructed to provide shelter; river ports are located on a river, and their waters are not retained by any artificial means. The right panel plots the estimated average port productivities, $\log \omega_{jt} = \log \omega_j + \delta t + \nu_{jt}$, as obtained from the estimation of the service time CES specification (8).

	Cobb Douglas	CES
α_{cd}	0.88	
β_{cd}	0.60	
α_{ces}		0.76
η		-0.16
N	263	263

Table 8: This table reports the estimated parameters from the service time CES specification (8), as well as a Cobb-Douglas specification, where α_{cd} is the exponent on L_{jt}/s_{jt} and β_{cd} is the exponent on c_{jt}/s_{jt}.

	γ_t	$\hat{\beta} y_t$
	(1)	(2)
Constant	-2.643	0.1122
	(6.185)	(0.0671)
coal price	0.0115	
	(0.0081)	
food price	-0.1428	
	(0.0488)	
agricultural price index	0.1912	
	(0.0675)	
base metal price	0.0080	
	(0.0170)	
BDI	0.0007	
	(0.0005)	
Industrial price index	-0.0277	
	(0.0665)	
$\hat{\beta} y_{t-1}$		0.9715
		(0.0131)
N	70	373
R^2	0.36	0.94

Table 9: This table reports the estimated parameters β that allow us to extrapolate to a longer time-series in the first column, as well as the estimated coefficients for the AR(1) process for γ_t given in (7) in the second column.

	Change in Wait Time (%)	Change in Trade (%)	Change in Welfare (%)	Rate of Return from Investment (%)
Alliance	0.00	0.00	0.00	0.00
Ama	-0.17	0.46	0.56	112.67
Baltimore	-0.12	0.85	0.60	67.54
Baton Rouge	-0.13	0.33	0.35	61.17
Beaumont	-0.11	0.28	0.33	103.02
Brownsville	-0.16	0.39	0.42	176.63
Burns Harbour	-0.00	0.00	0.00	0.89
Burnside	-0.23	0.31	0.34	84.75
Camden	-0.09	0.17	0.17	9.34
Charleston	-0.15	0.58	0.42	42.29
Cleveland	-0.02	0.03	0.07	16.91
Conneaut	-0.00	0.00	0.00	0.86
Convent	-0.12	0.60	0.55	251.83
Corpus Christi	-0.12	0.61	0.46	181.52
Darrow	-0.13	0.26	0.30	122.41
Davant	-0.14	0.45	0.49	85.48
Destrehan	-0.20	0.70	0.76	298.29
Detroit	-0.11	0.18	0.27	29.46
Duluth	-0.06	0.20	0.30	59.17
East Chicago	-0.04	0.05	0.14	25.24
Fairless Hills	-0.09	0.15	0.15	14.61
Galveston	-0.17	0.49	0.46	58.17
Gramercy	-0.37	0.37	0.37	86.60
Grays Harbor	-0.17	0.60	0.27	46.32
Hampton Roads	-0.21	1.32	0.90	117.97
Houston	-0.11	0.69	0.59	138.16
Jacksonville	-0.16	0.47	0.45	49.41
Kalama	-0.17	1.47	0.86	221.25
La Place	-0.13	0.21	0.23	45.14
Lake Charles	-0.13	0.34	0.36	117.08
Long Beach	-0.16	1.04	0.65	38.82
Longview	-0.11	0.86	0.50	63.07
Los Angeles	-0.16	0.57	0.41	18.41
Milwaukee	-0.09	0.17	0.32	28.69
Mobile	-0.14	0.67	0.61	132.08
New Orleans	-0.13	0.72	0.59	71.29
Newark	-0.12	0.40	0.30	16.86
Pascagoula	-0.23	0.59	0.54	90.05
Point Comfort	-0.03	0.06	0.08	50.38
Port Arthur	-0.20	0.76	0.70	244.81
Portland	-0.11	0.86	0.49	20.98
Reserve	-0.09	0.32	0.34	116.04
Richmond (CA)	-0.16	0.65	0.45	27.13
Savannah	-0.16	0.61	0.49	42.97
Seattle	-0.14	0.37	0.26	5.95
St Bernard Port	-0.14	0.29	0.29	65.82
Tacoma	-0.20	0.74	0.44	37.98
Tampa	-0.14	0.66	0.59	44.95
Toledo	-0.07	0.25	0.38	106.15
Vancouver	-0.17	1.18	0.66	61.26
Wilmington	-0.14	0.32	0.27	27.45

Table 10: For each port in our sample, the table reports the change in average wait times, aggregate trade, and aggregate welfare, when expanding the port's capacity so it can handle one more vessel, as well as the rate of return from the investment (the cost is just recovered when the rate of return is equal to 100).

	Welfare Gains	
δ_{geo} (standardized)	13.19	
	(4.19)	
δ_{geo} of substitute ports (standardized)	14.31	
	(14.18)	
Queue (standardized)		5.02
		(2.48)
Queue of substitute ports (standardized)		0.75
		(4.88)
R^2	0.12	0.21
N	51	51

Table 11: Determinants of Welfare Gains: This table regresses welfare gains on both the attributes of the treated port and those of its 10 closest substitutes. We standardize the covariates to ease interpretation of the coefficients.

B Creation of Port Infrastructure Dataset

In this Appendix, we discuss the creation of the port infrastructure dataset.[45] Because ports seldom share information about their infrastructure, we use geolocation data and satellite imagery (provided by Google Maps and Google Earth Pro) to collect this information. Our approach consists of combining a number of different data sources (AXS Marine, Marine Traffic, Google Maps and Google Earth Pro), as well as a set of well-defined rules and extensive collaboration among data-collectors to achieve a reliable dataset on port infrastructure. We next describe the steps of our strategy. It is worth mentioning that ports are greatly heterogeneous (e.g. ocean vs. river) which makes each port unique in terms of data collection. This is also why current AI techniques may have difficulty collecting this information.

First, we obtain port polygons from AXS Marine. The polygons come in the form of Google Maps and highlight the outline of the port and its extent. This outline marks the geographic borders of the port. This is the area in which the data collectors need to identify different aspects of port infrastructure, including berths, storage, cranes, and silos.[46] Figure 21 provides examples of polygons, which illustrate the considerable geographic heterogeneity. In particular, it depicts the polygons for Duluth (MN) and Baltimore (MD) which are visible different: Duluth is a simple harbor covering a relatively small area. In contrast, Baltimore is much more complex, as it includes both rivers and other coastal areas and is very extensive, so that a large area needs to be checked for port infrastructure.

We cross-validate this information by using public information from Marine Traffic, which also provides locations for some of the berths. We add any missing terminals/berths to the port. In addition to giving us confidence on the port area, Marine Traffic contains information on the type of berth: we only want to include berths that handle cargo (general cargo, dry bulk, but also oil and containers) and thus remove berths that cater to fishing ships, cruise ships, military ships, and recreational vessels. We designate each berth to its purpose (dry bulk, oil, or containers).

In the next step, the data collectors measure several features of port infrastructure for every port: the length of every berth, the storage areas and silos, the number of cranes, and whether the port has railroad access, and/or a refinery nearby. Figure 22 provides one illustrative example from the port of Newark; on the image one can see clearly a ship parked at berth, the designated storage space next to it, as well as the cranes (un)loading the vessel. We next discuss each feature in turn.

[45]We are most grateful to Pedro Degiovanni, as well as a Ian Bourgin, Victoria Healey, Shai Hirschl, Virginia Jablonski, Ellie Kominiarek, Emily Taylor, Robert Yampanis, Jiayi Zhang for their careful and dedicated work in creating this dataset.

[46]It is worth noting that at least two different data collectors worked on each port, while each member of the team was trained by an existing member and was asked to re-encode a complex port that was previously completed as part of their training.

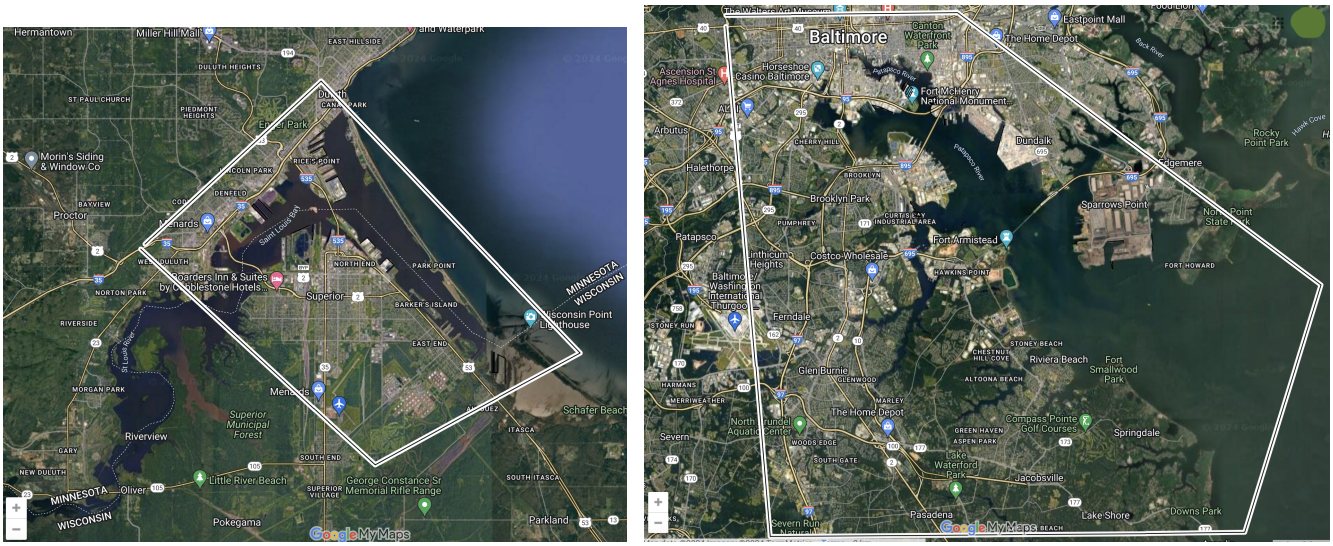

Figure 21: Port polygons for Duluth (MN) on the left and Baltimore (MD) on the right. These polygons capture the possible area of the port and are searched for port infrastructure, including berths, storage space, and cranes.

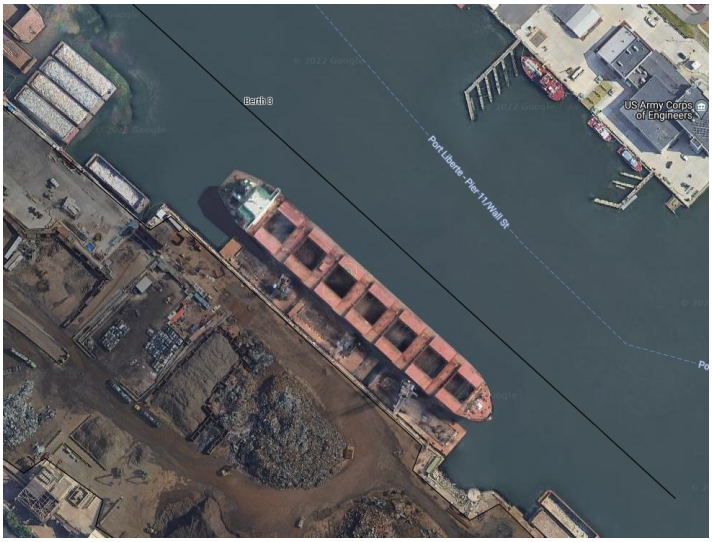

Figure 22: Satellite imagery of the port of Newark from Google Earth Pro. One can see a ship parked at berth, the designated storage space next to it, as well as the cranes (un)loading the vessel.

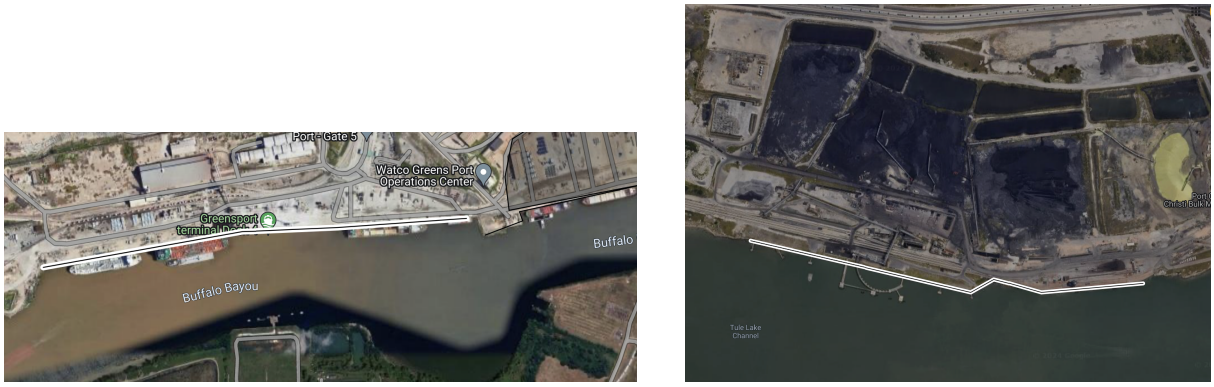

Figure 23: Different berths at the port of Houston (TX) on the left and Corpus Christi (TX) on the right.

Once all berths of a port are annotated in Google Maps, the data collectors locate them in Google Earth Pro. For each berth, they delineate its length and record it. Berths may have peculiar shapes so attention must be paid that all relevant berth space is included and that the measurement is done properly. As berths can have varying length, this measure of infrastructure is much more granular and informative than simply the number of berths. Figure 23 depicts examples of different berths from the ports of Houston (TX) and Corpus Christi (TX).

Ports have massive yards where they store cargo, such as bulk commodities, or containers. For each port, the data collectors first identify the different areas allocated to storage and for each of them they create a polygon on Google Maps. Then, they record its area in acres and indicate whether the storage is used for containers or not.

Figure 24: Different storage areas at Houston (TX) on the left and Hampton Roads (VA) on the right.

Encoding storage areas is likely the trickiest step of our approach and does in some cases require judgement calls by the data collectors. In those cases, they were asked to cross-check their work with

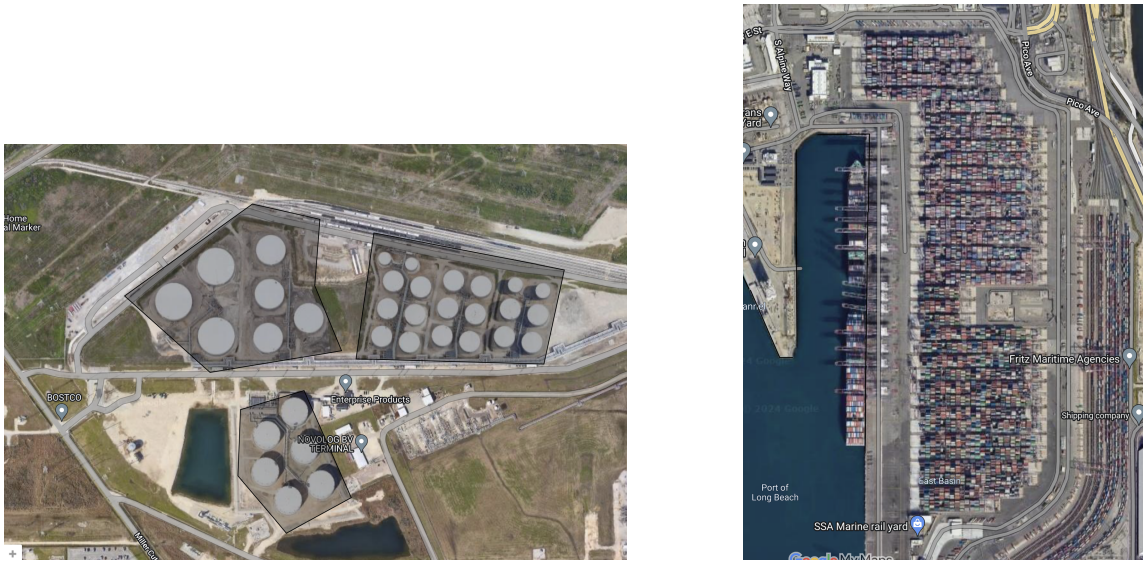

Figure 25: The left picture depicts silos areas at Houston (TX). The right picture depicts container storage space in Long Beach (CA).

their co-workers and mark their assessments in the dataset. For instance, coal is generally easy to recognize due to its color, and so are piles of metal, but other minerals may appear more like dirt (this is less of an issue in the US). Similarly, some storage sites may initially look like construction, but they are in fact deposits of minerals. Container storage is easier to identify. However, while some ports have large, distinct areas of container storage, other ports have more small, scattered areas, or areas that look like junkyards. Another issue involves distinguishing between storage of cars and parking lots.[47] To address the potential for measurement error when it comes to storage areas, we include in our analysis only storage areas that are within a one mile radius from a berth, and perform robustness by using storage areas within a three mile radius. Figure 24 depicts examples of storage areas in Houston (TX) and Hampton Roads (VA). The right picture in Figure 25 shows an example of container storage from Long Beach (CA).

Silos are often used to store grain, oil or gas and are an important part of storage for commodities. Data collectors count silos, which are giant cylinders that look like circles on Google Earth Pro and report their total area as part of storage.[48] The left picture in Figure 25 depicts silos in Houston (TX).

Next, we measure the number of cranes. It is easier to see cranes in sunnier places due to the shadows they cast; so oftentimes the data collectors have to search for a good recent satellite image on a sunny day to spot them. Note for instance the cranes in the port of Newark in Figure 22.

We also track whether the port is connected to a railroad or not, as well as whether there is a refinery

[47] In this case we adopted the following guidelines: if we see containers with a truck attached to them, we consider it a parking lot. If we see traffic signs and driving space between cars, we consider it a parking lot. If the cars are all packed together we count it as storage.

[48] To distinguish silos from other shorter cylindrical objects, we check whether they cast a shadow.

nearby. In the US, all ports in our sample are connected to railroad.

Finally, we use Google Earth Pro (in particular the "historical imagery" function) to construct a time series for all the features collected, from 2010 until 2021.

C Counterfactual Simulations

In this Appendix, we describe the algorithm employed to compute the equilibrium of our model. Every month t, the importer-exporter pairs face an aggregate demand state γ_t, while every port j offers a service rate $\mu_{jt}(= 1/T_{jt})$ and has capacity K_j.[49]

Finding the equilibrium of our model requires computing, at every port $j \in J$ and month $t \in \{1, \ldots, T\}$, the following three objects. First, the probabilities $\{p_{n,jt}(\tau_h)\}_{n,h}$ that there are n ships at port j at instant τ_h of month t, for a grid $\{\tau_1, \ldots, \tau_H\}$ of half-hour intervals within the month. Second, the average total time at port TT_{jt} implied by these probabilities. Third, the arrival rate λ_{jt} of ships at port j. Note that, since traders make their decisions based on the expected time at port TT_{jt}, the arrival rate λ_{jt} will depend on this object. In turn, the state of the queue at each port, described by the probabilities $\{p_{n,jt}(\tau_h)\}_{n,h}$ depends on the ships' arrival rates λ_{jt}. Our definition of equilibrium requires that the arrival rate λ_{jt} is consistent with TT_{jt} given our demand model, and that the probabilities $\{p_{n,jt}(\tau_h)\}_{n,h}$ are consistent with the arrival rate λ_{jt}, given the differential equations (3) describing the evolution of the queue. Also note that in every month the initial state of the queue depends on the residual queue at port at midnight on the last day of the previous month, which is described by the probabilities $\{p_{n,jt}(\tau_H)\}_n$. Beyond this "initial condition" the evolution of the queue is independent across months. We next describe how we compute the equilibrium objects within each month.

The algorithm proceeds as follows:

1. Make an initial guess for the expected time at port in month t, $TT_t^0 = \left(TT_{1,t}^0, \ldots, TT_{J,t}^0\right)$.

2. At each iteration k, we inherit TT_t^{k-1}. We then update our guess according to the following steps.

   (a) Compute the number of ships arriving at each port, $n_t^k = \left(n_{1t}^k, \ldots, n_{Jt}^k\right)$, given the expected

[49] In our simulations we set the service rate $\mu_{jt} = 1/T_{jt}$, where T_{jt} is the average service time in port j and month t. In practice, we use average service time over a six month period. Note that when a port invests in expanding its capacity K_j, the number of ships serviced contemporaneously can increase, thus driving up the average service time T_{jt} (see Equation 2). We thus assume that when increasing capacity, the port has to hire enough additional workers and buy enough additional cranes so as to keep the service time fixed. Therefore, in the simulations the service rate remains fixed at μ_{jt} when the port capacity increases. This is also a realistic assumption in our view: adding a new berth without manning it with workers and capital equipment may well not be reasonable.

time at port TT_t^{k-1}, from

$$n_{jt}^k = \sum_{fd} \frac{\exp\left(\delta_{jft} + \beta_d \text{dist}(j,d)\right)}{1 + \sum_l \exp\left(\delta_{lft} + \beta_d \text{dist}(j,d)\right)} M_{fd},$$

where

$$\delta_{jft} = \beta_f \text{dist}(f,j) - \beta_T TT_{jt}^{k-1} - \beta_p p_{jt} + \gamma_t + \gamma_f + \gamma_{l(j)} + \xi_{jft}. \tag{11}$$

Denote the corresponding arrival rates by $\lambda_{jt}^k = n_{jt}^k$.

(b) Next, compute the evolution of the queue in each port *within* month t, given the arrival rates λ_{jt}^k, service rates μ_{jt}, capacities K_j, and probabilities $p_{n,jt-1}(\tau_H)$. To this end, we consider a grid $\{\tau_1, \ldots, \tau_H\}$ of half-hour intervals within the month and apply the differential equations (3) to each port j as follows:

i. For half-hour $h = 1$ we approximate the probability $p_{n,jt}^k(\tau_1)$ that there are n ships at port j at instant τ_1 of month t using

$$p_{n,jt}^k(\tau_1) = p_{j,t-1,n}(\tau_H) + (\tau_h - \tau_{h-1}) \frac{d}{d\tau} p_{n,jt-1}(\tau_H)$$

where

$$\frac{d}{d\tau} p_{n,jt-1}^k(\tau_H) = \begin{cases} -\lambda_{jt}^k p_{0,jt-1}(\tau_H) + \mu_{jt} p_{0,jt-1}(\tau_H) & \text{if } n = 0 \\ -\left(\lambda_{jt}^k + n\mu_{jt}\right) p_{n,jt-1}(\tau_H) + \lambda_{jt}^k p_{n-1,jt-1}(\tau_H) + (n+1)\mu_{jt} p_{n+1,jt-1}(\tau_H), & \text{if } 0 < n < K_j \\ -\left(\lambda_{jt}^k + K_j\mu_{jt}\right) p_{n,jt-1}(\tau_H) + \lambda_{jt}^k p_{n-1,jt-1}(\tau_H) + K_j\mu_{jt} p_{n+1,jt-1}(\tau_H), & \text{if } n \geq K_j \end{cases} \tag{12}$$

ii. Similarly, for each interval $h > 1$, we approximate

$$p_{n,jt}^k(\tau_h) = p_{n,jt}^k(\tau_{h-1}) + (\tau_h - \tau_{h-1}) \frac{d}{d\tau} p_{n,jt}^k(\tau_{h-1}),$$

where $\frac{d}{d\tau} p_{n,jt}^k(\tau_{h-1})$ is defined as above.

(c) Given the probabilities $\left\{p_{n,jt}^k(\tau_h)\right\}_{h=1}^H$ we can compute the expected time at port as

$$TT_{jt}^k = \frac{1}{H} \sum_h \sum_n \left(\frac{1}{\mu_{jt}} + \frac{(n - K_j + 1)}{\mu_{jt} K_j} \mathbb{I}\{n \geq K_j\} \right) p_{n,jt}^k(\tau_h),$$

where we use the probabilities $p_{n,jt}^k(\tau_h)$ to take expectations across the possible states of the queue at port j, and then take the average across the intervals τ_h.

(d) If $\max_j \left\{ \left| TT_{jt}^k - TT_{jt}^{k-1} \right| \right\} < \epsilon$, then we set $TT_{jt} = TT_{jt}^k$ and $\lambda_{jt} = \lambda_{jt}^k$.

In practice, in our counterfactuals we compute the equilibrium across 300 simulations, where we re-draw the monthly demand state γ_t. For each simulation m we set γ_0^m equal to the estimated demand state on January 2018. Then, we simulate the demand state for thirty years according to

$$\gamma_t^m = c + \rho \gamma_{t-1}^m + \epsilon_t^m$$
$$\epsilon_t^m \sim N\left(0, \sigma_\epsilon^2\right),$$

where $(c, \rho, \sigma_\epsilon^2)$ are the estimated parameter of the AR process described in Section 4.2.

Finally, we compute expected discounted welfare by taking the average over the simulations. We compute the expected discounted welfare for each simulation as

$$\sum_t \beta^t W\left(TT_t, \gamma_t\right),$$

where

$$W\left(TT_t, \gamma_t\right) = \frac{1}{\beta_p}\left[\sum_{d,f} \log\left(\sum_j \exp\left(\beta_T TT_{jt} + \gamma_t + x_{jfd}\right) + 1\right) M_{df} + \gamma^{euler}\right]$$

i.e. the present discounted sum of consumer welfare across all options, $W\left(TT_t, \gamma_t\right)$, which depends on the vector of time at port TT_t, the aggregate demand γ_t, as well as x_{jfd} which includes the remaining port characteristics, such as the distances and prices. Given that port prices in our data vary little over time, even during the period when demand surged, we set them equal to their port-specific sample average. Similarly, we keep ξ_{jft} constant and equal to the average over time in our sample. We calibrate the discount factor, β, to 0.95 annually.

D Cost of Port Infrastructure

In this Appendix we provide all details for the construction of the cost of infrastructure from Section 6.1. We begin with the cost of dredging and then move to the cost of land.[50]

Dredging Costs As discussed in Section 6.1, there are two main steps to estimating dredging costs. For each port j, we need (i) a port- or berth- level measure of the seafloor depth at which rock material

[50]The analysis in this section would not have been possible without the invaluable research assistance of Zach Saunders.

begins, denoted d_j^r; (ii) an estimate of the function linking dredging costs, to the quantity of rock and non-rock material extracted. Once the function from (ii) is known, we use the amount of rock and non-rock material that needs to be extracted at each port to compute port-specific dredging costs. We describe each step in turn.

For the first step we rely on the National Crustal Model (NCM), currently in development by the US Geological Survey (USGS). This uses a large set of existing geophysical data – including measures of porosity, lithology, compressional and shear wave velocities, and depth to bedrock and basement material – to estimate an array of geophysical properties on a 1-km by 1-km grid spanning the entire contiguous United States. For each berth in our dataset, we use unpublished meter-by-meter estimates of subsurface compressional (i.e., primary) wave velocities estimated by the NCM that the USGS has provided. The model structure, calibration, and estimation procedure used to derive these estimates are developed in Boyd (2020).[51] From these meter-by-meter estimates, we calculate the depth (in meters) below each berth at port j at which primary wave velocity first exceeds 2,300 meters per second. This threshold is used as the point of delineation between sediment (i.e., non-rock) and rock material, whereafter drilling and blasting methods must be used to prepare subsurface material for dredging. The choice of 2,300 m/s primary wave velocity to delineate rock from non-rock material follows previous literature (see, e.g., MurphyIII et al. (2011) and sources therein) and is informed by measured ripping productivity of subsurface materials at different primary wave velocities.[52]

We now turn to the second step: estimating the cost of dredging as a function of the quantity of sediment that is excavated and depending on its consistency (rock vs. non-rock). Over the last two decades, the US Army Corps of Engineers (USACE) has conducted several large-scale dredging projects at major US harbors. For two such projects – the Delaware River Main Channel Deepening Project and the Boston Harbor Deep Draft Navigation Improvement Project, which broke ground in 2010 and 2017 respectively – the USACE provides detailed estimates of predicted dredging costs. For each proposed design depth under consideration, estimates were produced for the quantity of rock and non-rock material to be dredged, total dredging and disposal-related costs both for rock and for non-rock materials, total

[51] We are most grateful to Oliver Boyd for making this data available to us.

[52] Generally, most rock types transition from 'rippable' to 'marginal' between 2,000 and 3,000 m/s, whereas the transition to 'non-rippable' generally occurs between 3,000 and 4,000 m/s. While these measures might suggest a threshold of 3,000 m/s be used as the point at which drilling and blasting become necessary for the preparation of rock for excavation, we note that measured ripping productivity (in cubic meters of ripped material per hour) declines precipitously well before rock is classified as either marginal or completely non-rippable. MurphyIII et al. (2011) correlate measured ripping productivity of a D-10 tractor ripper against measured compressional wave velocity for a number of sandstone ripping projects and show that measured productivity can fall very close to 0 m^3/h at seismic velocities of less than 2,500 m/s, well below the 'marginal' classification threshold for sandstone rock. Similar results can be found in productivity ranges reported in the 49th edition of the Caterpillar Performance Handbook.

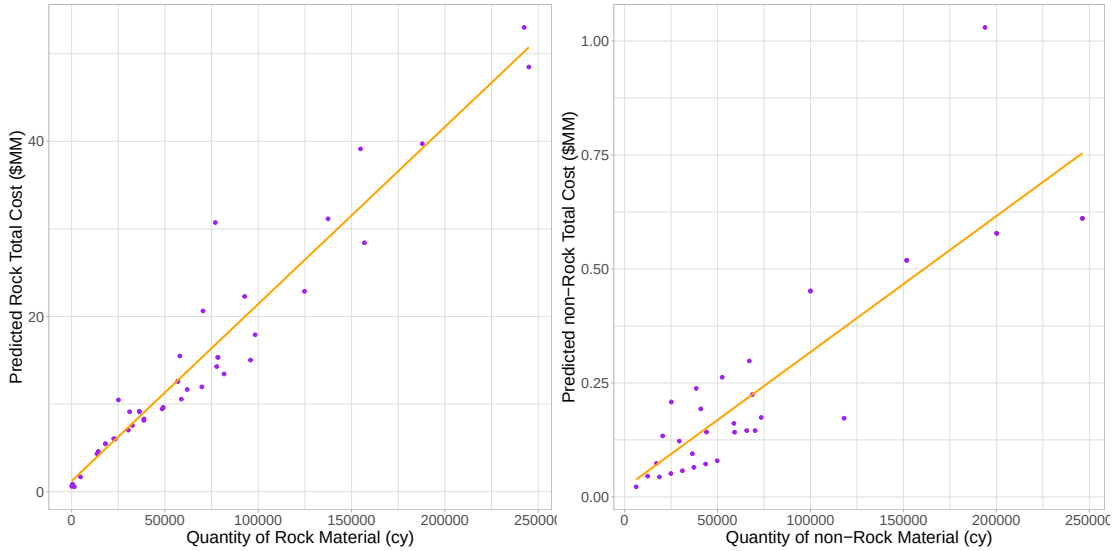

Figure 26: Dredging costs for rock and non-rock material. The dots correspond to different quantity-cost pairs for each project-berth-depth reported by the USACE and the solid line represents the estimated functional form of the dredging cost function.

drilling and blasting-related costs required to prepare rock material for dredging and disposal, and the price level at which these dollar costs were calculated. These estimates were provided not just at the project level, but also for each terminal berth and for each major section of harbor channel. Altogether, there are 78 distinct port-section-depth combinations for which the USACE provides dredging cost and material quantity estimates.[53]

After separating these rock and non-rock dredging costs, all additional costs – including items such as "Miscellaneous Costs," "Construction Contingency," "Planning Engineering and Design," "Supervision & Administration," "Construction Management," "Resetting of Aids to Navigation," and "Real Estate" – are apportioned as rock- and non-rock-related costs in proportion to the ratio of rock to non-rock-related dredging construction costs calculated above. Finally, all prices are converted to 2022 USD by adjusting for inflation using the PPI.

We then regress linearly the USACE costs on the quantity of excavated material, in order to obtain the dredging cost functions, $C_j^{d,r}(q^r), C_j^{d,nr}(q^{nr})$. As independent variables we use the quantities of rock and non-rock material that need to be dredged for each project-segment-depth reported by the USACE, while the dependent variable is the associated cost estimated by the USACE as required to remove the quantity of the material. The fit for this specification is quite good, as shown in Figure 26.

[53]For the Delaware River deepening project, all data come directly from the 2002 Delaware River Main Channel Deepening Project Comprehensive Economic Reanalysis Report, Appendix A: Cost Estimate. For the Boston Harbor deepening project, all cost and quantity estimates are retrieved from the 2013 Boston Harbor Navigation Improvement Study Final Feasibility Report, Appendix D2 – Project Cost Estimates.

Finally, we use this function to extrapolate to all ports and obtain a measure of dredging there. The total dredging cost is equal to the cost of dredging rock plus the cost of dredging non-rock material:

$$C_j^d(l,w,d) = C_j^{d,r}(q^r) + C_j^{d,nr}(q^{nr})$$

where q^r (q^{nr}) is the quantity of rock (non-rock) material. The task now is to compute q^r and q^{nr} for any port. We have

$$q^r = \mathbb{I}\left\{\left(d_j^r \leq d+2\right)\bigcap\left(d_j < d+4\right)\right\} lw\left(\left(d+4\right) - \max\left\{d_j, d_j^r\right\}\right)$$

and

$$q^{nr} = \max\left\{0, lw\left(\min\left\{d_j^r, d+2\right\} - d_j\right)\right\}$$

where d_j is the current depth of port j, d_j^r is the depth at which rock begins at port j, d is the desired depth at which digging will take place, and l, w, d is the length, width and depth/draft of the average ship respectively. Let us explain these expressions:

First, suppose that $d_j \leq d+2 \leq d_j^r$, i.e., the new berth can be dredged to depth $d+2$ (to service ships of depth d with the customary 2 ft of additional overdepth) without hitting any non-dredgeable rock material. In this case, no rock material is dredged, $q^r = 0$, while the quantity of non-rock dredged material is given by $q^{nr} = lw\left(\left(d+2\right) - d_j\right)$.

Second, suppose that $d_j^r \leq d_j \leq d+2$. In this case, all material to be excavated will consist of non-dredgeable rock material; i.e., $q^{nr} = 0$ and $q^r = lw\left(\left(d+4\right) - d_j\right)$. Given that the bottom of the berth will now consist of exposed rock, the customary 4 ft of overdepth must be provided to prevent damage to ships' hulls resulting from accidental contact with the rocky berth floor.

Third, suppose $d_j \leq d_j^r \leq d+2$. Then, non-rock material will be dredged from the actual depth d_j to the depth of hitting rock d_j^r, while non-dredgeable rock material must be removed from depth d_j^r to depth $d+4$, in which case $q^n = lw\left(d_j^r - d_j\right)$ and $q^r = lw\left(\left(d+4\right) - d_j^r\right)$.

For the purposes of this exercise, we set d_j equal to zero ft, as we consider this a brand new berth; we have however performed robustness with 5 ft and 10 ft, as a minimum depth below the water level; in these cases, costs are lower and a few more ports exhibit positive net returns. Moreover, the length of the berth is given by parameter κ_1 of the Leontief capacity function (9) as described in Section 4.1. We set the width (depth) equal to the width and draft of the average vessel.

Land Costs The second term of the cost function for investment in infrastructure consists of land costs required for additional storage. As discussed in Section 6.1, we estimate the cost of land acquisition by first using open-source map data to identify parcels of industrial and commercial land near existing port infrastructure, and then we use highly granular land value estimates from the ecology and conservation literature to assign unit land prices to these parcels.

In particular, we use OpenStreetMap (OSM) to extract all land-based features with either an industrial or commercial use located within 1, 2, or 3 miles of any berth affiliated with each port in our data.[54] The distance used for each port is the minimum distance sufficient to capture at least one parcel of industrial or commercial land neighboring a local terminal.[55]

We next calculate the monetary value of this industrial land that we have identified as local for each port. To do so, we rely on Nolte (2020) who provides land values for the US. In particular, Nolte (2020) uses data on land sales to train a predictive model for expected sale prices for all private properties in the contiguous United States.[56] The resulting property-level land price estimates are then rasterized into a 6,054 by 9,618 element raster covering the landmass of the contiguous US. Thus, each resulting raster element corresponds to an approximately square parcel of land with side length measuring 480 meters and is assigned a single value measuring its estimated price-per-hectare.

Using all such parcels of industrial and commercial land that are local to the ports in our dataset, we calculate the price-per-hectare of local industrial and commercial land as the weighted average of rasterized land value estimates from Nolte (2020), with weights corresponding to the amount of industrial and/or commercial land falling within each square of the land value grid. Because the land values are measured in 2010 USD, we then adjust for inflation using the PPI.

Finally, we compute the total land costs as the product of land prices and the acreage required to handle the new vessel, which is given by parameter κ_2 of the Leontief capacity function (9) as described

[54] OSM is an open-access online map service updated and maintained by its users and comprises a highly-detailed database of natural and man-made geographic features. Among other information, OSM includes shape files representing the outline and location of each constituent geographic feature along with a land-use variable indicating the primary use of each land-based feature. The land-use variable takes 80 possible values, including education, residential, commercial, and industrial, among others.

[55] In two cases – the Port of Los Angeles and the Port of Long Beach – a single geographic feature must be dropped following this procedure, as it corresponds to a large area of Los Angeles Harbor that has been incorrectly coded as industrial land in OpenStreetMap.

[56] In particular, Nolte (2020) uses a tree-based bagging algorithm to estimate the expected sale price of all private land in the contiguous United States, with the unit of observation being the tax-assessor parcel. The model is trained on 6 million private land sales taking place between 2000 and 2019, and estimated land price is modeled as a function of publicly-available data on land ownership, sale information, building footprints, terrain, accessibility, land cover, hydrology, flood risk, demographics, and environmental protections. The trained model and underlying data are then used to predict expected sale prices for all private properties in the contiguous United States, expressed in 2010 USD. These estimates are validated against verified sale prices of 4,128 publicly funded land acquisitions in 659 US counties, and found to explain approximately 70% of variation in log per-area sale prices.

in Section 4.1.

Bulkhead Costs A bulkhead is a vertical shoreline stabilization structure that primarily retains soil, and provides minimal protection from waves. As already mentioned, data on bulkhead construction are rare, however we are able to use reported information from the Charleston (SC) Harbor port expansion project. As part of this project, the construction of a new bulkhead was necessary, which – at the project planning stage – had an associated estimated cost of $22 million measured at April 2014 prices. The terminal has 3,800 linear ft of berthing space, which results in a rough cost estimate of $5,789 per linear ft of bulkhead. After adjusting for inflation, we let bulkhead costs amount to $6,600 per linear foot of berth in 2022 USD. This estimate is in line with one more reported bulkhead construction project we were able to find, the 1,671-ft Port Fourchon bulkhead construction project completed by the Greater Lafourche Port Commission in 2019, which cost approximately $11 million, corresponding to a cost-per-foot of roughly $6,582 in 2022 USD.

References

Alessandria, G., S. Y. Khan, A. Khederlarian, C. Mix, and K. J. Ruhl (2022): "The Aggregage Effects of Global and Local Supply Chain Disruptions: 2020-2022," *mimeo, University of Rochester*.

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

Almagro, M., F. Barbieri, J. C. Castillo, N. G. Hickok, and T. Salz (2024): "Optimal Urban Transportation Policy: Evidence from Chicago," Tech. rep., National Bureau of Economic Research.

Anderson, J. E. and E. Van Wincoop (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93, 170–192.

Arellano, M. and S. Bond (1991): "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *The Review of Economic Studies*, 58, 277–297.

Arnott, R., A. De Palma, and R. Lindsey (1990): "Economics of a bottleneck," *Journal of urban economics*, 27, 111–130.

Asturias, J. (2020): "Endogenous Transportation Costs," *European Economic Review*, 123, 103366.

Bai, X., J. Fernández-Villaverde, Y. Li, and F. Zanetti (2024): "The causal effects of global supply chain disruptions on macroeconomic outcomes: evidence and theory," *NBER Working Paper*, 32098.

Bailey, S. (2021): "Competition and Coordination in Infrastructure: Port Authorities' Response to the Panama Canal Expansion," *Unpublished Manuscript*.

Berry, S., M. Carnall, and P. Spiller (2006): "Airline Hubbing, Costs and Demand," Elsevier, vol. 1 of *Advances in Airline Economics, Vol. 1: Competition Policy and Anti-Trust*, 183–214.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 841–890.

Bhat, U. N. (2008): *An Introduction to Queueing Theory: Modeling and Analysis in Applications*, vol. 36, Springer.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Bordeu, O. (2023): "Commuting Infrastructure in Fragmented Cities," *Job Market Paper, University of Chicago Booth School of Business*.

Boyd, O. (2020): "Calibration of the U.S. Geological Survey National Crustal Model," Tech. rep., US Geological Survey, USA.

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

Brancaccio, G., M. Kalouptsidi, T. Papageorgiou, and N. Rosaia (2023): "Search Frictions and Efficiency in Decentralized Transport Markets," *Quarterly Journal of Economics*, 138, 2451–2503.

Brooks, L., N. Gendron-Carrier, and G. Rua (2021): "The Local Impact of Containerization," *Journal of Urban Economics*, 126.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021): "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *The Quarterly Journal of Economics*, 136, 1255–1321.

Donaldson, D. (2012): "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *forthcoming, American Economic Review*.

Donaldson, D. and R. Hornbeck (2016): "Railroads and American Economic Growth: A "Market Access" Approach," *Quarterly Journal of Economics*, 131, 799–858.

Downs, A. (1962): "The Law of Peak-Hour Expressway Congestion," *Traffic Quarterly*, 16, 393–409.

——— (1992): *Stuck in traffic: Coping with peak-hour traffic congestion*, Brookings Institution Press.

Ducruet, C., R. Juhasz, D. K. Nagy, and C. Steinwender (2022): "All Aboard: The Effects of Port Development," *mimeo, Columbia University*.

Duranton, G. and M. A. Turner (2011): "The fundamental law of road congestion: Evidence from US cities," *American Economic Review*, 101, 2616–2652.

Durrmeyer, I. and N. Martinez (2022): "The welfare consequences of urban traffic regulations," *TSE Working Paper*.

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

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

Friedt, F. L. and W. W. Wilson (2020): "The Economic Impact of Seaport and Other Infrastructure Investments and Leakages: A Literature Review," *Journal of Transport Economics and Policy*, 25, 219–243.

Fuchs, S. and W. F. Wong (2022): "Multimodal Transport Networks," *mimeo, University of Oregon*.

Ganapati, S., W. F. Wong, and O. Ziv (2022): "Entrepot: Hubs, Scale, and Trade Costs," *mimeo, Georgetown University*, forthcoming.

Glaeser, E. L. and J. M. Poterba (2021a): *Economic Analysis and Infrastructure Investment*, Chicago: University of Chicago Press.

——— (2021b): "Economic Perspectives on Infrastructure Investment," in *Rebuilding the Post-Pandemic Economy*, ed. by M. S. Kearney and A. Ganz, Aspen Institute Press.

Grieco, P. L., C. Murry, J. Pinkse, and S. Sagl (2022): "Conformant and efficient estimation of discrete choice demand models," *Unpublished Manuscript*.

Grossman, G. M., E. Helpman, and H. Lhuillier (2021): "Supply Chain Resilience: Should Policy Promote Diversification or Reshoring?" *NBER Working Paper*, 29330.

Heiland, I., A. Moxnes, K. H. Ulltveit-Moe, and Y. Zi (2021): "Trade From Space: Shipping Networks and The Global Implications of Local Shocks," *mimeo, University of Olso*.

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

Hummels, D., V. Lugovskyy, and S. Skiba (2009): "The Trade Reducing Effects of Market Power in International Shipping," *Journal of Development Economics*, 89, 84–97.

Hummels, D. and S. Skiba (2004): "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture," *Journal of Political Economy*, 112, 1384–1402.

Ishikawa, J. and N. Tarui (2018): "Backfiring with Backhaul Problems: Trade and Industrial Policies with Endogenous Transport Costs," *Journal of International Economics*, 111, 81–98.

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

Kreindler, G. (2023): "Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications," *NBER Working Paper*, 30903.

Larson, R. C. and A. R. Odoni (2007): *Urban Operations Research*, Dynamic Ideas, 2nd ed.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.

MurphyIII, W., W. Ward, B. Boyd, W. M. IV, R. Nolen-Hoeksema, R., M. Art, and D. Rosales-R (2011): "Geophysical, Geological, Geotechnical, and Mechanical Testing of Rock." *Proceedings of the Western Dredging Association (WEDA XXXI) Technical Conference and Texas AM University (TAMU 42) Dredging Seminar*, 423–435.

Nolte, C. (2020): "High-resolution land value maps reveal underestimation of conservation costs in the United States," *Proceedings of the National Academy of Sciences*, 117, 29577–29583.

Notteboom, T., A. Pallis, and J.-P. Rodrigue (2022): *Port economics, management and policy*, Routledge.

Redding, S. J. (2016): "Goods trade, factor mobility and welfare," *Journal of International Economics*, 101, 148–167.

Redding, S. J. and M. A. Turner (2015): "Chapter 20 - Transportation Costs and the Spatial Organization of Economic Activity," in *Handbook of Regional and Urban Economics*, ed. by G. Duranton, J. V. Henderson, and W. C. Strange, Elsevier, vol. 5 of *Handbook of Regional and Urban Economics*, 1339–1398.

Sytsma, T. and W. W. Wilson (2020): "Port Choice and International Trade in Agricultural Products," *mimeo, University of Oregon*.

Vickrey, W. S. (1969): "Congestion Theory and Transport Investment," *The American Economic Review*, 59, 251–260.

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