Transportation Bottlenecks and Aggregate Shocks

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

New York University, Harvard University, Boston College Harvard University

September 2026

Do transportation disruptions affect aggregate outcomes? We develop a tractable model of the transportation sector to study its impact on prices and international trade. By casting maritime transportation as a queueing network, the model identifies where the transportation bottlenecks lie and delivers a simple and easy-to-compute measure of transportation capacity over time. We calibrate the model using shipping and trade data and use it to evaluate the resilience of the shipping sector during the last 20 years. Port bottlenecks increased commodity prices and led to shortages during the Great Supply Chain Disruption, although a contemporaneous negative supply shock mitigated these effects, while insufficient shipping capacity played a key role during the 2000s Commodity Supercycle.

Keywords: transportation, shipping, ports, bottlenecks, disruptions, supply chains, capacity, prices, queueing network, resilience

[*]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 332, Chestnut Hill, MA, 02467, USA, theodore.papageorgiou@bc.edu

[§]Harvard University, Department of Economics, Littauer Center, Cambridge MA 02138, USA, yixinzhou@g.harvard.edu

[¶]We gratefully acknowledge financial assistance from NSF Award SES-231533-231534-231535. We are thankful for several helpful comments from George Alessandria, Costas Arkolakis, Amit Khandelwal, and Jean-Paul L'Huillier.

1 Introduction

A barrage of disruptions related to transportation has affected the global economy in recent years: from the Great Supply Chain Disruption following the COVID-19 pandemic, the Red Sea crisis and the tensions in the Strait of Hormuz, to the environmental shocks in critical waterways, the transportation sector that has seemingly carried international trade seamlessly during the era of globalization, has proven to be a critical bottleneck. With all this happening amid large and persistent increases in prices, as well as goods shortages, many have speculated about the impact of transportation disruptions on aggregate economic activity. This paper tackles precisely this question. How much do transportation disruptions matter in the aggregate? And how does the transportation sector affect the economy's response to (aggregate) shocks?

We focus on maritime transportation (ships and ports), which is responsible for more than 80% of total trade in goods and featured prominently in the public debate surrounding the Great Supply Chain Disruption.[1] We develop a tractable model of the transportation sector to study its impact on prices and international trade. The key ingredient is a queueing network representation of the transportation sector, which pins down where bottlenecks arise and delivers a simple, easy-to-compute measure of capacity over time. We calibrate the model using shipping and trade data and use it to evaluate the resilience of transportation during the last 20 years. Port bottlenecks increased commodity prices and led to shortages during the Great Supply Chain Disruption, although a contemporaneous negative supply shock mitigated these effects, while insufficient shipping capacity played a key role during the 2000s Commodity Supercycle.

More specifically, we begin by presenting descriptive evidence that the transportation sector has limited ability to adjust its capacity in periods of increased demand. We focus on dry bulk shipping, which transports raw materials such as coal, grain, and iron ore. We show that the two key components of the sector, i.e. port infrastructure and the shipping fleet, are rigid. Port capacity expands only sluggishly, while fleet size responds little to demand shocks, partly because of long shipbuilding lags (Kalouptsidi, 2014, Dunn and Leibovici, 2025). Existing ships also operate at persistently high utilization rates. Next, we examine whether transportation disruptions and bottlenecks pass through to commodity prices. To get a first descriptive pass at this question, we estimate the relationship between commodity prices and port congestion using a novel instrumental variables strategy based on unexpected disruptions in port

[1] At the same time, maritime transportation accounts for a substantial share of sourcing lead times: of the roughly 55 days required to deliver production inputs, more than half are spent at sea or in ports on average (Alessandria et al., 2023).

service times. We find that port congestion explains three percentage points of the increase in commodity prices in 2022, or about 12% of the total increase.

We then embed a model of the transportation sector into a simple trade model to explore the impact of these rigidities on commodity prices and trade. Trade flows from an exporter to an importer and must be executed by the transportation sector, which can create bottlenecks that limit the quantity of goods that can be traded. We model the transportation sector as a sequence of queues: exporters need to first wait for a ship, then they must wait to enter and get serviced at port, and finally they must travel to the destination. Using insights from queuing theory, we construct a measure of the transportation system's capacity that is simple and easy-to-compute. We show that this measure depends on the existing stock of port infrastructure, the number of ships, as well as the port service time and distance traveled. This measure is useful beyond our setting: because it is computed from widely available data, it can offer researchers and policymakers a transparent, high-frequency gauge of transportation bottlenecks that can be tracked in real time.

More specifically, the transportation sector reaches capacity when any of the queues in the network reaches capacity. As a result, bottlenecks arise either because of port capacity constraints or a shortage of ships. Port capacity, i.e. the number of shipments processed per period, is determined by port infrastructure and port service times. Thus, even with an abundant fleet, port capacity limits throughput. Conversely, ship capacity depends on the fleet size, but also on port infrastructure, service times, and travel distances. Longer port stays or greater distances keep ships occupied for longer, reducing the effective fleet available to transport goods.

Our setup is general enough to accommodate different types of disruptions, including port service time disruptions (e.g. strikes, COVID-19 inspections), port closures (e.g. Baltimore bridge collapse, closure of Hormuz), fleet disruptions (e.g. fines to Chinese built ships in the US), disruptions to distance travelled (e.g. closure of the Suez canal, or Houthi attacks), as well as more complicated disruptions including interactions between such events and aggregate demand or supply shocks, such as the Great Supply Chain Disruption. Our model thus allows us to decompose and evaluate the impact of each.

In equilibrium, when none of the transportation constraints bind, shipping prices equal the marginal cost of shipping; however when one of the constraints binds, shipping prices jump up so that excess demand equals the available supply of transportation services, clearing the market. This two-tier structure of prices implies that prices are often roughly constant at relatively low levels, but when the constraint binds, they can increase substantially, depending on the excess demand for shipping services. This sharp implication

of the model is in fact corroborated in the data: shipping prices, are fairly stable most of the time, but then rise sharply and eventually revert. These spikes in shipping prices are significantly higher than what would be implied by any increase in the underlying marginal costs (mostly fuel costs). During the last 20 years the sector experienced two such episodes: one during the Great Supply Chain Disruption following the COVID-19 pandemic, and one during the 2000s Commodity Supercycle.

We next pull together data from various sources to calibrate our setup. Our empirical exercise focuses on global ports and on dry bulk shipping. Our main data involves the universe of port calls, which inform us on service times and congestion at all world ports, as well as port handling capacity. In addition, we use shipping data, such as shipping prices, the global fleet, fuel costs, etc. We combine this with bilateral trade data (quantities and prices).

We use the model to construct a measure of transportation capacity during the last 20 years, from 2004 to 2023. We compare capacity to the quantity transported during this period and find that capacity was indeed binding during much of the 2004-2009 period, as well as from 2021 to 2023. In other words, exactly during the Great Supply Chain Disruption and the 2000s Commodity Supercycle. Importantly, in the early period capacity was determined by ship bottlenecks, whereas in the latter one the binding constraint was on the port side.

How much did these bottlenecks affect commodity prices and trade? We begin with the Great Supply Chain Disruption, which was a combination of a positive demand shock, a negative supply shock, and a transportation disruption, primarily in port service time. During 2021–2023, COVID-19-related labor restrictions increased port service times, while demand for raw materials rose sharply. At the same time, supply contracted because of lockdowns and, later, the war in Ukraine. The increase in port service time reduced the capacity of the transportation sector, which combined with the increase in demand led to the transportation constraint binding. However the reduction in supply made the role of the transportation during this period less pronounced, as the reduction in supply reduced the need for transportation services. In a counterfactual world with slack transportation capacity, we find that commodity price inflation would have been three percentage points lower which corresponds to 13% of the total increase. Instead, a counterfactual world without the negative supply shock would have led to equally high prices, with the transportation sector now responsible for 72% of the increase.

The 2000s Commodity Supercycle was driven by a positive demand shock stemming from China's surging demand for raw materials, which was met by insufficient shipping capacity. The ship bottleneck became binding and shipping prices soared. Counterfactual simulations suggest that the transportation

sector was responsible for up to 23% of the increase in commodity prices during this period. Unlike the latter period, now it was the ship constraint that was binding. Following this period there was a large increase in shipbuilding orders and the global dry bulk fleet has more than doubled.

The impact of disruptions depends on the extent of slack in the system. As an example, we examine the Houthi attacks in the Red Sea, which caused ships to reroute around the Cape of Good Hope. Despite the substantial increase in voyage distances, the effect on shipping prices was relatively limited because the ship constraint for bulk shipping was not binding at the time. As a result, the decline in the effective supply of ships caused by the rerouting had little effect on shipping prices in dry bulk shipping, other than through higher marginal (fuel) cost.

Our analysis illustrates that maritime chokepoints are not just local logistics risks. They are a channel through which geopolitical, climate, and infrastructure disruptions can become shocks for the entire transport system, by reducing global transportation capacity. They also suggest that differences in how congestion rents are captured across ships and ports may have important implications for investment incentives and, ultimately, the evolution of transportation capacity; we discuss this in more detail in the Conclusion.

Related Literature

This paper relates to a rapidly growing literature studying disruptions—including supply chain frictions, transportation bottlenecks, and geopolitical shocks—and their effects on aggregate outcomes. Much of this literature was motivated by the disruptions surrounding the COVID-19 pandemic, although its focus has broadened to encompass a wider range of shocks affecting production, transportation, and international trade. See for instance, Carvalho, Nirei, Saito, and Tahbaz-Salehi [2021], Grossman, Helpman, and Lhuillier [2021], Elliott, Golub, and Leduc [2022], Alessandria, Khan, Khederlarian, Mix, and Ruhl [2023], Calvo Lopez and Golub [2023], Comin, Johnson, and Jones [2023], Amiti, Itskhoki, and Weinstein [2024], Castro-Vincenzi [2024], Comin, Johnson, and Jones [2024], Bernanke and Blanchard [2025], Gopinath, Gourinchas, Presbitero, and Topalova [2025], Bai, Fernández-Villaverde, Li, and Zanetti [2026b], Bai, Fernández-Villaverde, Li, Marto, and Zanetti [2026a], Blaum, Esposito, and Heise [2026], Fuchs, Leibovici, and Wong [2026], Kleinman, Liu, Redding, and Xu [2026], Kohler, L'Huillier, Phelan, and Weiss [2026], Korovkin, Makarin, and Miyauchi [2026].

A closely related literature studies transportation disruptions more directly. Carrière-Swallow, Deb, Furceri, Jiménez, and Ostry [2022] and Michail, Melas, and Cleanthous [2022] explore the role of freight

rates (the Baltic Dry Index, BDI) on inflation; Benigno, di Giovanni, Groen, and Noble [2022] construct a measure of transportation and supply chain disruptions (the global supply chain pressure index) and explore its impact on inflation. In recent work, Bai, Fernández-Villaverde, Li, and Zanetti [2024] explore the impact of port congestion on inflation in the US and Kanzig and Raghavan [2026] also examine the impact of canal closures on prices and other macro outcomes. All these papers, using for the most part regressions or VAR analysis, find that shipping disruptions have a sizable impact on the economy. Comin, Johnson, and Jones [2024] provide a stylized model to study the role of import constraints, transportation being one of them. Dunn and Leibovici [2025] provide a model for ship investment and utilization to study their impact on macro outcomes. In a related paper, Brancaccio, Kalouptsidi, and Papageorgiou [2024], we examine the policy implications of disruptions on port infrastructure and we compute the optimal level of port investment.

Our main contribution is to develop a tractable equilibrium framework of the transportation sector that explicitly models both fleet capacity and port infrastructure. The framework can be used to quantify how disruptions propagate through transportation, distinguish transportation bottlenecks from other shocks affecting prices and trade, and evaluate the effects of a broad range of transportation disruptions. Moreover, it offers a real-time measure of the capacity of the transportation sector from widely available data so that policymakers can track the extent of slack in the system. Unlike the existing literature, most of which treats transportation disruptions as exogenous shocks, our focus is on modeling the transportation sector and the equilibrium role of transportation capacity in shaping the transmission of disruptions and aggregate shocks. This allows us to provide a clear mechanism for the formation of shipping prices, whose extreme volatility has proven difficult to capture in most existing models.

More broadly, our paper also contributes to an expanding literature in economics studying maritime transportation, e.g. Hummels and Skiba [2004], Hummels [2007], Hummels, Lugovskyy, and Skiba [2009], Kalouptsidi [2014], Ishikawa and Tarui [2018], Asturias [2020], Bailey [2021], Brooks, Gendron-Carrier, and Rua [2021], Wong [2022], Brancaccio, Kalouptsidi, Papageorgiou, and Rosaia [2023], Ducruet, Juhasz, Nagy, and Steinwender [2024], Ganapati, Wong, and Ziv [2024], Brancaccio, Kalouptsidi, and Papageorgiou [2025], Heiland, Moxnes, Ulltveit-Moe, and Zi [2025], Fuchs and Wong [2026].

Finally, this paper naturally relates to the large literature that has focused on the impact of transportation infrastructure on trade (e.g. Donaldson, 2012, Allen and Arkolakis, 2014, Redding and Turner, 2015, Donaldson and Hornbeck, 2016, Redding, 2016, Brancaccio, Kalouptsidi, and Papageorgiou, 2020, Fajgelbaum and Schaal, 2020, Allen and Arkolakis, 2022, Redding, 2025), as well as the international

trade literature studying 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 presents the data, and documents key empirical patterns. Section 3 introduces our framework for the transportation sector. Section 4 presents our model calibration and results. Section 5 quantifies the impact of the transportation sector on aggregate outcomes during the last two decades. Section 6 concludes.[2]

2 Data and Empirical Patterns

In this section, we describe the data that we use and then present some key empirical patterns highlighting the rigidities in maritime transportation and their impact.

2.1 Data

We focus on trade in commodities and raw materials, such as grain, coal, iron ore, and fertilizers, which is executed by dry bulk carriers. 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 (Brancaccio et al., 2020). Dry bulk shipping accounts for roughly one-third of global seaborne trade by volume (UNCTAD), while dry bulk carriers account for roughly 43% of the world fleet's carrying capacity, which includes also containerships and oil tankers.[3]

We collect information on (i) detailed ship and port flows; (ii) freight rates, as well as the worldwide dry bulk fleet; (iii) global trade quantities and values. We describe each briefly and provide further details in Appendix B.

First, we use the universe of port calls between 2013 and 2023 Q2 for bulk carriers larger than 10,000 DWT. The dataset is obtained from AXS Marine (AXS Dry). This dataset tracks the different stages a ship goes through during a port call: 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. Each

[2] Appendix A provides additional figures and tables, Appendix B provides details on the data and the descriptive analysis. Appendix C includes the proofs and Appendix D reports details on the calibration.

[3] Bulk ships are very different from containerships. While dry bulk shipping (and tankers) are an unconcentrated industry, consisting of a large number of small ship-owning firms that operate like taxis, 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 (as would oil tankers). Indeed, container shipping features 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 hold for both containers and oil tankers.

observation consists of 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 is loaded or unloaded, an estimate for the commodity on board, and the size of the vessel. 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.

Our port call dataset consists of about 2.1 million port calls between 2013 and 2023, involving 1853 ports in 144 countries. By construction this dataset provides service times (i.e. time spent at berth for loading or discharge) separately from wait times in the queue. Ships spend on average 5.7 days at port; of this, 55% (3.2 days) consists of service time, while 44% (2.5 days) of wait time at anchorage (i.e. time spent waiting to get serviced at congested ports). We focus on the following major commodities: coal, grains (wheat, soybean, corn), iron ores, fertilizers, and cement.

Second, we obtain transport prices from a dataset of dry bulk shipping contracts, collected by Clarksons Research. The dataset is a sample of voyage trip-charter contracts for 2002-2015 and 2020-2023. An observation is a transaction between a shipowner and a charterer for a specific trip. We observe the vessel, the charterer, the dates, the price in dollars per ton for each trip, as well as information on the origin and destination. We rely on this dataset to compute a measure of bilateral transport prices for our sample period. Our final sample consists of 21,352 voyage trip-charter contracts covering 66 countries, with an average (median) freight rate of 16.3 (12.5) USD per ton.

Additionally, we use a monthly time series of the global fleet from Clarksons Research, which reports the world's total stock of vessels measured both by number and by deadweight in every month from 2000 to 2023. For fuel costs, we rely on a time series of bunker prices collected by Refinitiv.

To obtain exporter prices for commodities we rely on the Harmonized System (HS) bilateral trade data from the United Nations COMTRADE dataset, which is available at a monthly frequency from 2010 to 2023 and annual frequency from 2002 to 2009. As this dataset reports both quantity and value of bilateral trade, we use the ratio of the two to obtain the bilateral exporter price. This yields a mean (median) exporter price of 535.5 (264.4) USD per ton across commodities. Importer price, instead, is measured as the sum of exporter price plus the transport prices described above.[4]

[4]Following Amiti et al. [2024], we use exporter-reported value on a free-on-board basis (FOB) to measure prices, rather than importer-reported value, because certain importers do not consistently report trade value of imports in CIF.

2.2 Transportation Sector Capacity and Adaptability

We next examine the shipping sector's capacity and whether it has the ability to adjust in periods of high demand for shipping services, such as during the Great Supply Chain Disruption. The two key determinants of the shipping sector's capacity are (i) port infrastructure and (ii) the fleet of ships. We examine each in turn.

Ports

Can port capacity adjust to absorb demand during booms? As a first piece of evidence, the long queues and severe port congestion in the post COVID-19 period that made numerous headlines, is suggestive of the rigidity of port infrastructure.

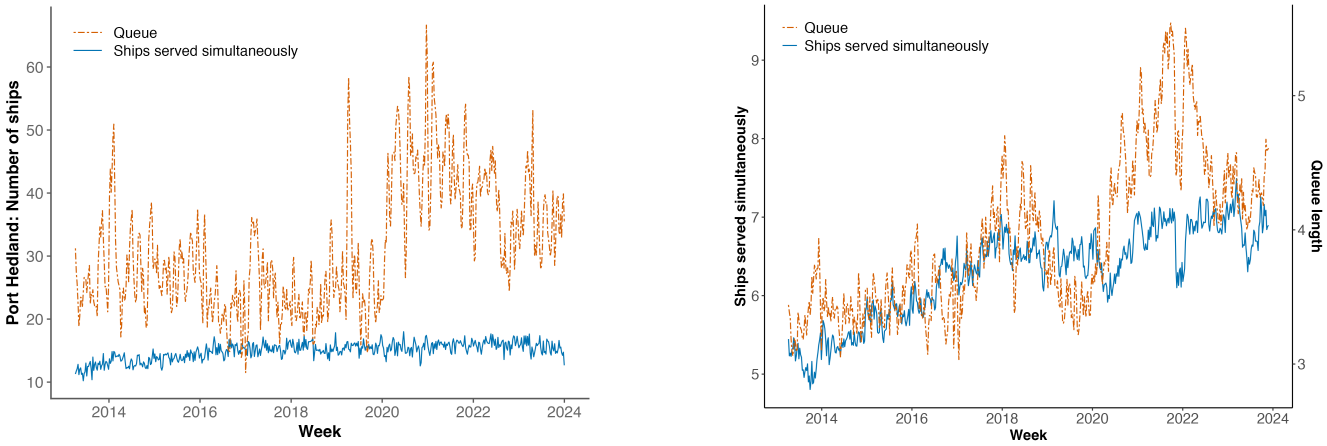


Figure 1: The left panel plots the number of ships simultaneously serviced and number of ships in queue at the Port of Hedland, Australia. The right panel plots the number of ships simultaneously serviced (left axis) and number of ships in queue (right axis) for the average port.

To assess whether port capacity can absorb demand shocks, we examine how many ships a port services simultaneously. If ports cannot increase the number of ships that they simultaneously service, they experience queue buildup after a shock. As an example, the left panel of Figure 1 plots the number of ships serviced by the port of Hedland (a major Australian port for iron ore exports), as well as the average queue length, capturing demand. Notably, it is the queue rather than the throughput that reflects the fluctuations in demand; indeed, the number of ships serviced remains remarkably flat over time, suggesting port Hedland is at capacity throughout this period.

The right panel of of Figure 1, generalizes this across ports by plotting the average number of ships that ports service simultaneously, as well as the average port queue length, which captures port demand. Up until approximately 2018, the number of ships that ports service slowly rises, largely in line with long-

run increases in demand (since, unlike port Hedland, most ports are not always operating at capacity). After that point however, the number of ships serviced simultaneously is largely flat, seeming to have hit a plateau, while the queue increases sharply, suggesting that ports have reached the limit of their capacity.

Port capacity is a fairly complex object and depends on physical infrastructure such as the number of berths, storage areas, connectivity, etc, as well as employment, capital equipment and port productivity (see Brancaccio et al., 2024). The evidence here illustrates that indeed these factors, many of which are sluggish to adjust, become binding constraints when demand rises.

Ships

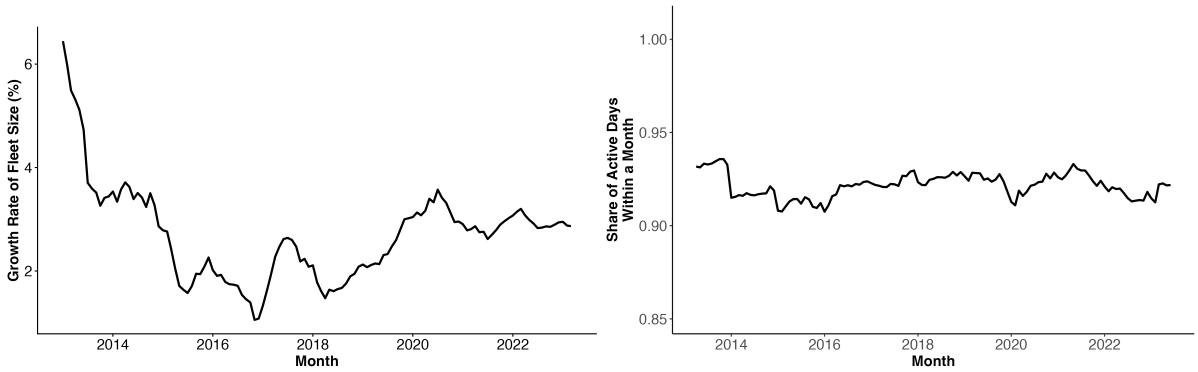


Figure 2: Left panel: Annual growth rate of fleet size. Right panel: Fraction of ships that are active every month. Active ships are either traveling loaded or empty (ballast) or they are waiting for a shipment.

We next turn to ships and consider whether the fleet can respond to demand shocks or whether fleet capacity, too, is fixed in the short-run. On the extensive margin, Kalouptsidi [2014] documents that time-to-build in ship production, typically ranging from two to five years, prevents the fleet from expanding rapidly during periods of high demand (see also Dunn and Leibovici, 2025 for the case of containerships). In addition, even without time-to-build, given that ships are long-term investments (average lifespan of a ship is over 30 years), it is not obvious that ship construction would be responding, if the demand shock was perceived to be short-lived. The left panel of Figure 2 confirms this pattern in our sample period: the growth rate of the fleet is fairly low, at about 3% annually. The initially higher growth rate in the early 2010s is further evidence of the mismatch due to time to build, as it captures delivery of ships that were ordered prior to the Great Recession.

On the intensive margin, the right panel of Figure 2 plots the average fraction of days within a month that a ship spends *active* at sea, i.e., traveling loaded or empty, waiting at port, or for shipment.[5] Ships

[5]Inactive ships include those in shipyards, as well as those that are taken out of active service and "laid up" . These classifications are based on the AXS data, which report the status of each vessel on a daily basis.

on average spend more than 90% of the time within a month active, and this share varies little cyclically, suggesting that ships are not being "mothballed" when demand is low. This is partly because there are substantial costs to inactivity (mainly capital deterioration, as humidity and seawater corrode the vessel and marine growth accumulates on the hull). As a result, there does not appear to be a buffer of inactive ships that can be deployed when demand picks up.

2.3 Port Disruptions and Prices

Do these bottlenecks, due to rigid port and ship capacity, passthrough to commodity prices? A recent growing literature (e.g. Benigno et al., 2022, Carrière-Swallow et al., 2022, Michail et al., 2022, Bai et al., 2024, and Kanzig and Raghavan, 2026) has engaged with this question, by looking at whether different measures of transportation disruptions have a significant impact on prices and/or inflation. Several proxies for disruptions have been studied, such as freight rates, delivery lags, inventory stocks, but much focus has also been given to port congestion. This is partly motivated by the correlation between port congestion and commodity prices that both rose dramatically post COVID-19, as shown in Figure 3, and discussed in numerous press headlines.

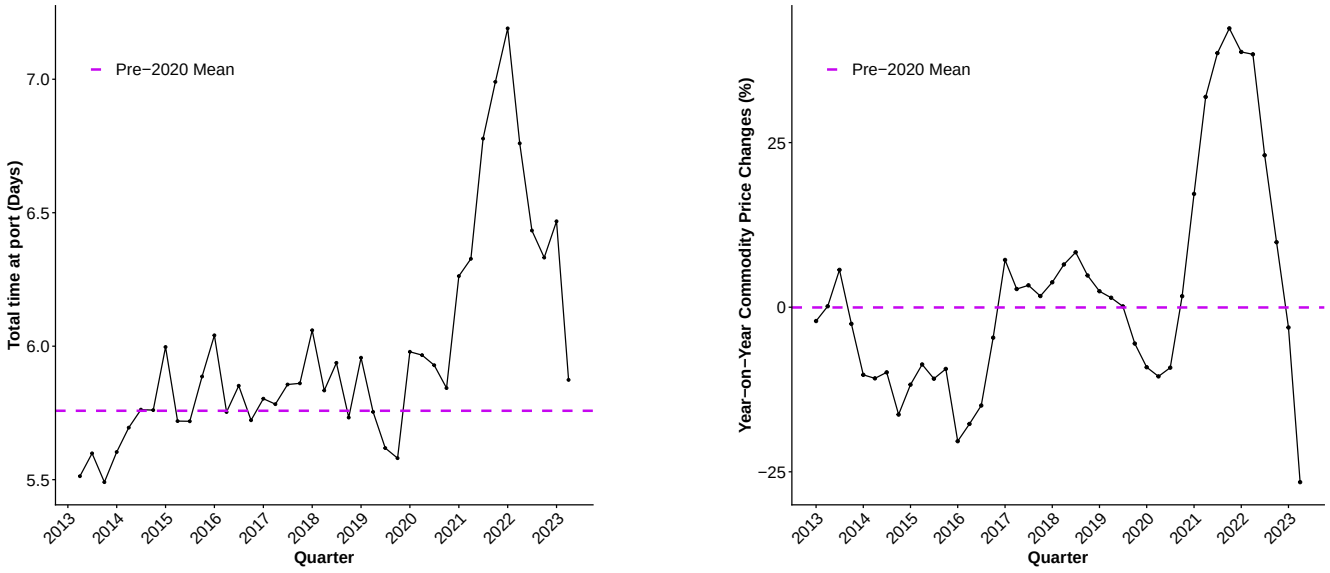


Figure 3: Port Congestion and Commodity Price Changes. The left panel depicts average total time at port across ports measured in days. The right panel depicts quarter-to-quarter percent change in commodity prices, where prices are calculated as described in Appendix B.

As a first pass at answering this question, we estimate the impact of port congestion on commodity prices within a regression framework (much like the literature above); we then return to this question using our model in order to decompose the impact of different disruptions and shocks, as well as gain

insights on the mechanisms at play.

We test whether total time spent at a loading port has a positive impact on commodity prices, by estimating the following specification,

$$\log\left(P_{ijct}/P_{ijc,t-4}\right) = \alpha_{ij} + \gamma_{it} + \lambda_{jt} + \delta_{ct} + \beta TT_{ijct} + \epsilon_{ijct}, \tag{1}$$

where t is a quarter, i, j is a pair of importing county i and exporting country j, and c is the commodity. The dependent variable is the year-over-year log change in P_{ijct}, the export price for commodity c from j to i, in quarter t. The key regressor TT_{ijct} is the time ships spend at ports in the producing country j during quarter t. Specifically, TT_{ijct} is a weighted average of port-level time at port, across the ports in producing country j through which commodity c ships to importer i, with weights given by each port's share of that trade. The terms α_{ij}, γ_{it}, and λ_{jt} are fixed effects. It is worth noting that trade in these commodities is highly concentrated among a small number of exporting countries, whereas demand for imports is much more fragmented. For example, Australia alone accounts for 54% of world iron ore exports; Indonesia and Australia together account for about 60% of world coal trade; and Brazil accounts for 51% of world soybean trade. This concentration is important because it means congestion at a single exporter's ports can meaningfully affect trade and prices, with few alternative origins available to absorb the disruption.

Estimating equation 1 directly suffers from a serious endogeneity problem: time spent at the port is correlated with unobserved commodity demand or supply shocks. For instance, an increase in the demand for commodities would lead to higher commodity prices, but would also increase the number of ships arriving at port and, therefore, increase time at port through increased congestion. In fact, congestion mechanically increases when demand rises. Thus, the OLS estimate of the coefficient of interest β would be upward biased. Similarly, a reduction in the supply of commodities can lead to a (spurious) negative correlation between prices and time at port. A supply shortage, driven for instance by closures in production sites, would increase commodity prices; but it would also lead mechanically to lower port congestion. Hence we may obtain a downward bias on the coefficient of interest, β. An instrument is therefore necessary.

We construct an instrument for TT_{ijct} as follows. The total a time a ship spends at a port is the sum of the service time, i.e. the time it takes to discharge or load cargo, plus the time it spends outside the port waiting to be serviced (i.e. in congestion). We instrument total time at port by using unexpected

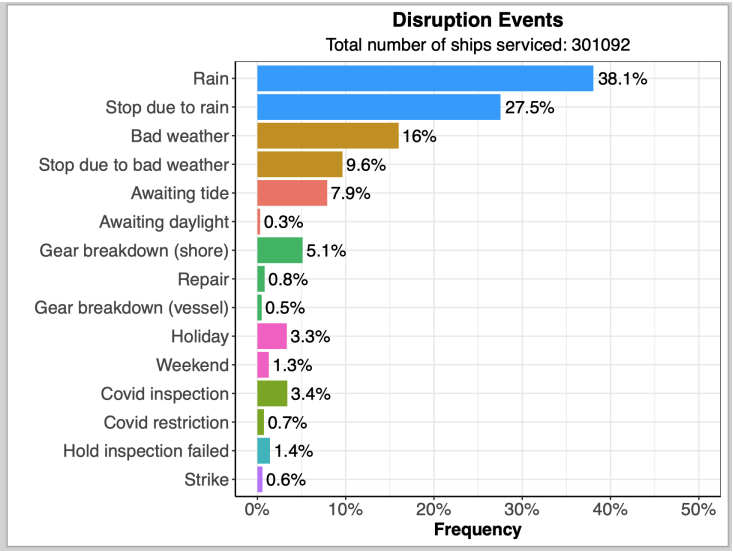


Figure 4: Port Disruption Events. This figure reports the percent of port calls that were affected by each of these disruptions. Disruption categories are not mutually exclusive. The data is constructed from port logs that report every single event that occurs during a port call.

disruptions in *service time* that are plausibly independent of supply and demand shocks.

What are such disruptions? To better understand them, we explore a dataset provided to us by a (by request, anonymous) company that collects data on port charges and port events. The data are essentially detailed port logs, i.e. detailed listings of all events that occur during a port call, such as cargo handling, pilotage, and inspections. From these, we isolate events that form disruptions that have stopped or delayed service. Figure 4 illustrates that disruptions are quite frequent and often attributed to weather events, as well as gear breakdowns.

To capture these types of disruptions, we construct an instrument by flexibly regressing service times on a number of controls and taking the residual. Appendix B provides the details, as well as the first stage results (see Table 3 there).[6] Moreover, Table 4 in the same appendix provides a validation of our instrument by correlating it with the most frequent disruptions reported in Figure 4. For instance, it shows that the instrument is correlated with average precipitation within a 10 km radius of the port during each port call, even after controlling for rain in the broader area.

Column (2) in Table 1 reports the IV estimate of the impact of port congestion on price increases. The estimates suggest that an additional day at port leads to approximately a 1.4 percentage point increase in commodity prices. For comparison, Column (1) reports OLS estimates of (1). We find that

[6]In addition, the impact of such unexpected disruptions on time at port may be nonlinear. For example, large disruptions such as labor strikes may have a larger impact on time at port than a small operational disruption (Brancaccio et al., 2024). We therefore create an alternative instrument as a robustness, where we interact the unexpected disruption with a dummy indicating whether the shock is large (above the 75th percentile by importer-exporter-commodity pair). We find that results are robust.

the OLS estimates are substantially lower than the IV estimates, suggesting that supply shocks introduce a downward bias that overall dominates the bias in the OLS estimates. This pattern is not entirely surprising and is corroborated by our estimated model later on; it is driven both by COVID-19 related closures in production sites, but also by the war in Ukraine, which severely reduced the worldwide supply of grain and fertilizers over our sample period.[7]

	OLS	IV
Model:	(1)	(2)
Constant	-0.0153**	
	(0.0062)	
TT at port	0.0089***	0.0140**
	(0.0010)	(0.0060)
importer-exporter		Yes
commodity-quarter		Yes
importer-quarter		Yes
exporter-quarter		Yes
R^2	0.001	0.22
Observations	61,946	60,983
KP F Stat, TT at port		150.0

Table 1: Impact of time at port on commodity price increases. The table reports the results from estimating equation (1). The dependent variable is log four-quarter commodity price change. The independent variables are total time at port, importer-exporter fixed effects, commodity-quarter fixed effects, importer-quarter fixed effects and exporter-quarter fixed effects. Column (1) reports results from OLS. Column (2) reports results from the IV regression using the instrument of unexpected service time disruptions. Standard errors are clustered at the importer–exporter level and reported in parentheses; *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 5 in Appendix B reports the estimates for the impact of port congestion on total trade in commodities. In particular, we estimate specification (1) with the year-on-year (log) change in the total quantity traded between two countries i and j, in commodity c. The IV estimates in Column (2) suggest that a one-day increase in total time at port leads to a 4.5 percentage point drop in the growth rate of quantity traded.

What do these estimates imply for the increase in commodity prices during the Great Supply Chain Disruption? Time at port increased by almost two days between 2019 and the first semester of 2022,

[7]Note also that the OLS might suffer from the classical attenuation bias, as total time at port may be a noisy measure of port disruptions. Our instrument of unexpected disruptions in service time could correct for this type of measurement error.

the peak of price inflation. This increase implies a 2.5 percentage point increase in the first half of 2022, which corresponds to about 12% of total commodity price inflation during this period.

3 A Model for Trade and Transportation

The empirical patterns documented in the previous section paint a picture of a sector that has very limited ability for adjustment in periods of high demand. In what follows, we embed a model of the transportation sector into a simple trade model to explore the impact of these rigidities on commodity prices and trade, as well as to identify where the transportation bottlenecks lie. The model delivers a simple and easy-to-compute measure of the transportation sector's capacity over time. This measure is useful beyond our setting: because it can be calculated from widely available data, it can offer researchers and policymakers a transparent, high-frequency gauge of transportation resilience that can be tracked in real time.

3.1 Environment

There are two countries, an exporter and an importer country and a single good to be traded. Trade must be executed by the transportation sector. For notational simplicity, we suppress time subscripts. We focus on the stationary equilibrium.[8]

Transportation Sector

The global shipping fleet and port infrastructure are the two key elements of the transportation sector. We describe each in turn.

There are S homogeneous ships in the world and each can transport one unit of the good.[9] Each ship is owned by an absentee shipowner. Travel time between the importer and exporter country is stochastic and a ship traveling from one to the other arrives at Poisson rate, with expected travel time d each way.

The expected cost of each round-trip is given by,

$$c(d) = c^0 + 2c^f d$$

[8] Since we focus on the trade in commodities, one-directional trade is consistent with the large trade imbalances observed in the data (Brancaccio et al., 2020).

[9] It has been documented extensively that the assumption on ship homogeneity is a reasonable approximation, as there is very little vessel price, freight rate or routing heterogeneity across ships (see Kalouptsidi 2014, Brancaccio et al. 2020). Similarly, for the assumption on constant returns to scale with respect to a shipowner's fleet (so that a ship is a firm).

where $c^0, c^f > 0$. Here c^f captures the fuel cost, so that $2c^f d$ is the total expected fuel consumption of the round-trip, and c^0 is the trip's fixed cost.

On the port side, ports in the exporting country can handle up to K ships at a time. If a ship arrives at port and there are fewer than K ships currently in service, the port is operating under capacity, and the ship is able to start loading its cargo immediately. Otherwise, the ship has to wait in a queue. Port service times are exponentially distributed, with expected service time denoted by T. Service time at the importing port is assumed to be zero.

Transportation as a Queueing Network

We next characterize the capacity of the transportation sector, i.e., the maximum sustainable throughput it can achieve when operating at full load. To do so, we treat the transportation sector as a sequence of queues (or a queuing network), that cargo needs to go through. The queueing network's capacity is then given by the minimum of the capacities of each queue. We rely on queueing theory to measure these objects.

Exporters arrive at the exporting port at Poisson rate q, which we later endogenize. First, the exporter may have to wait for a ship to arrive (the Ship Queue). Then once paired with a ship, the cargo proceeds to the port (the Port Queue) where it starts loading immediately if there are fewer than K ships currently being serviced, or it has to wait for a berth to open up. Finally, once loading is completed, the exporter's cargo travels to the importing port and arrives at rate $1/d$ (the Travel "Queue").

Formally, a queue, can be defined by three components: an arrival rate, a service rate, and a number of servers; this is represented in the form A/B/c, where A and B indicate the probability distribution of user inter-arrival times and service times, respectively, and c is the number of identical parallel servers. For instance, in the standard M/M/c queueing model, the inter-arrival and service times are exponential, while in the G/G/c queueing model, the distributions are general. These three parameters determine the stability of the queue, as well as other key objects, such as the expected total time in the system (Larson and Odoni, 2007).[10]

[10] As an example, consider the textbook M/M/c queueing model. Denote the Poisson arrival rate by λ. Denote the exponential service rate by μ. The expected total time in the queueing system is equal to the sum of the mean service time $1/\mu$, and the average time spent at queue, denoted by W, which is given by the well-known Erlang formula, $W = \frac{C(c,\lambda/\mu)}{c\mu-\lambda}$, where $C(c,\lambda/\mu) = \dfrac{\dfrac{(\lambda/\mu)^c}{c!}\dfrac{1}{1-\rho}}{\displaystyle\sum_{n=0}^{c-1}\frac{(\lambda/\mu)^n}{n!} + \frac{(\lambda/\mu)^c}{c!}\frac{1}{1-\rho}}$, and $\rho = \frac{\lambda}{c\mu}$. Thus expected total time in the system is given by $\frac{1}{\mu} + W$. Note that the wait time W is increasing in the arrival rate λ, as more arriving agents lead to longer queues and wait times; it is increasing in the expected service time, $\frac{1}{\mu}$, as longer service times make congestion more likely; and finally it is decreasing in

We next describe each queue in turn. The queueing network is shown schematically in Figure 5.

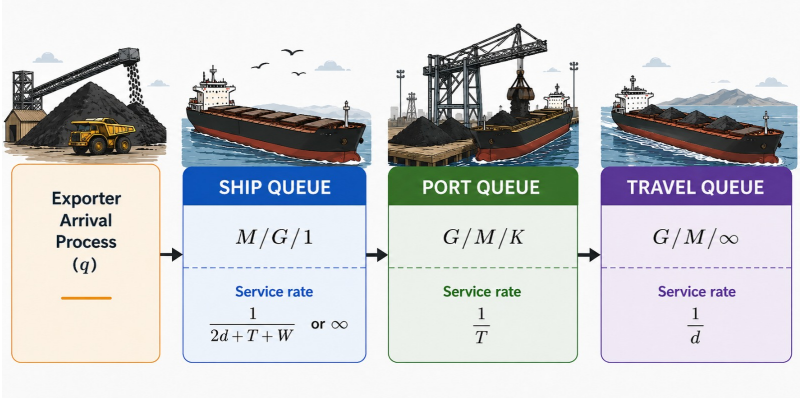


Figure 5: Transportation as a Sequence of Queues. Exporters arrive at rate q. They enter the M/G/1 Ship Queue, whose service rate is either infinite (if a ship is already waiting at port), or equal to $1/(2d+T+W)$, where W here denotes the time spent waiting in port congestion (W_{port}), plus the time (if any) ships have to wait in order to find an exporter. Once serviced at the Ship Queue (i.e. once they find a ship), exporters enter the G/M/K Port Queue, whose service rate is $1/T$. Finally, once serviced at the Port Queue they enter the G/M/∞ Travel Queue, whose service rate is $1/d$ and once serviced, they exit the system.

We begin with the **Ship Queue**. This queue can be modeled as a state-dependent M/G/1 queue: exporters arrive at Poisson rate q (hence the letter M). They get "serviced" once a ship arrives (hence the one server).[11] What is the service time? If upon arrival there are no available ships, the exporter needs to wait, and the service time is given by the arrival rate of ships, to which we return below. If instead an exporter arrives and finds a ship already waiting, then they instantaneously move to the Port Queue. Formally, this means the realized service rate is infinite in that case, so that service is instantaneous.

Second, the **Port Queue** can be modeled as a state dependent G/M/K queue. The arrival rate coincides with the departure rate from the Ship Queue above, the service rate is Poisson with mean $1/T$, and there are K servers. Note that the process of the arrival rate is general rather than Poisson (hence the letter G), since the output process from the Ship Queue is not necessarily Poisson.

The expected total time spent at the Port Queue is the sum of the expected service time, T, plus the wait time in port congestion, W_{port}. The wait time depends on the Port Queue's service time T and servers K, as well as the arrival rate at the port (see footnote 10). The latter is the departure rate of the Ship Queue, which is pinned down by the Ship Queue's arrival rate q and its service rate; the service rate, in turn, depends on the rate at which ships complete loops, as we show below. Note that

c, as more servers alleviate congestion and reduce expected wait times. Outside the M/M/c model, closed-form expressions for the wait time, W, are rare. However W can be either simulated or approximated by results in the queueing theory literature (see e.g. Whitt, 1983).
Finally it is worth pointing out that in the M/M/c model, the departure rate is Poisson. This property does not hold in other queues.

[11] To see why there is one server, note that the first exporter in line is in service and waiting for a ship to arrive (i.e. for service to complete), while any other exporters behind form a queue waiting to begin service.

$W_{\text{port}}\ (q, S, d, K, T)$does not have a closed-form expression, but can be simulated.[12]

Finally, the **Travel Queue**, which essentially is the travel time, can still be modeled as a queue, and in particular as a state dependent G/M/∞ queue with exponential service time with mean d, and infinite servers (thus, in that "queue" one only stays for their service time, there is never any wait time and the queue is just a delay factor). As before, the arrival process coincides with the departure process from the Port Queue and is thus not Poisson (hence the letter G).

Given the model parameters, this sequence of queues determines the time it takes for a cargo to be delivered at destination, and the stationary quantity of goods that is delivered (i.e. traded) within each period, which we denote by $Q\left(q, S, d, K, T\right)$.

Transportation Sector Capacity

Measuring the queueing network's capacity amounts to measuring the capacity of each queue described above. Following queueing theory, the key parameter governing stability of a queue is the utilization rate, $\rho \equiv \frac{\lambda}{\mu c}$, where λ is the arrival rate, μ is the service rate, and c is the number of servers. As long as $\rho < 1$, the departure rate coincides with the arrival rate, that is, the system is stable and reaches a steady state in which throughput equals the input rate and the server keeps up with arrivals on average. Instead, a queue reaches capacity when $\rho = 1$, corresponding to 100% utilization. In this case, the queue grows without bound and departures saturate at μc. Throughout, we require $\rho < 1$, for each queue. In what follows we measure the capacity of each queue, by considering throughput as its ρ approaches 1.[13]

We discuss each queue in turn, starting from the end.

The G/M/∞ Travel Queue is never at capacity since there are infinite servers, and therefore $\rho = 0$ for this queue. It thus plays no bottleneck role in this sequence of queues. Note, however, the Travel Queue adds time to the delivery of the goods, and therefore it affects the total quantity that can be delivered within each period.

Let us turn to the G/M/K Port Queue. This queue approaches full utilization, as its $\rho = \lambda / \left(K/T\right)$ approaches 1, where λ is the arrival rate for this queue. Denote by $\bar{Q}_{\text{port}}$ the maximum port throughput

[12] We assume that $W_{\text{port}}\left(q, S, d, K, T\right)$ is continuous and (weakly) increasing in q. This is intuitive, since the wait time should increase as demand rises. It can also be shown to hold if wait time is approximated using standard methods in the queueing theory literature (e.g. Whitt, 1983).

[13] For state-dependent queues, this condition is understood asymptotically: utilization must remain below 1 as the number of agents in the system grows.

per unit of time, while ensuring that the queue remains stable ($\rho < 1$), i.e.

$$\bar{Q}_{\text{port}} = \sup_{\lambda} \left\{ \lambda \in \mathbb{R}^{+} : \lambda < \frac{K}{T} \right\}$$

or

$$\bar{Q}_{\text{port}} = \frac{K}{T}$$

Intuitively, if $K = 1$, since the average service time per ship is T, the port can service up to $1/T$ ships on average per period. In the more general case, the port can handle up to $K\frac{1}{T}$ ships on average.

Finally, we turn to the Ship Queue. By definition the queue approaches full capacity when the rate of arrivals of exporters is high relative to how fast they are serviced– meaning there is always an exporter waiting. We can therefore focus on the case where ships have no idle time and every exporter must wait for service.

What "services" an exporter is a ship showing up to carry their goods. So the service rate of the Ship Queue is the rate at which ships become available to exporters, i.e., the ship arrival rate. Since ships loop between the two ports, this arrival frequency is directly determined by the average round-trip duration, from when a ship first arrives at the exporting port until it comes back to it:

$$2d + T + W_{\text{port}}\left(q, S, d, K, T\right)$$

In words, the ship needs to travel from the exporting to the importing port and back, which on average takes $2d$ periods; it spends $T + W_{\text{port}}\left(q, S, d, K, T\right)$ at the exporting port, where T is the expected port service time and $W_{\text{port}}\left(q, S, d, K, T\right)$ is the expected time spent waiting because of port congestion. Hence, each ship arrives at the exporting port at an average rate of

$$\frac{1}{2d + T + W_{\text{port}}\left(q, S, d, K, T\right)}$$

and since there are S ships in the fleet, the average rate at which a ship arrives at the exporting port, i.e. the service rate of the Ship Queue, becomes

$$\frac{S}{2d + T + W_{\text{port}}\left(q, S, d, K, T\right)} \tag{2}$$

Note that port congestion enters the cycle duration through $W_{\text{port}}\left(q, S, d, K, T\right)$ because the time a

ship spends waiting in the Port Queue is time it cannot return to pick up the next exporter. A longer $W_{\text{port}}(q, S, d, K, T)$ thus raises the trip duration and lowers the ship arrival rate, tightening the effective supply of ships available to exporters.

Denote by $\bar{Q}_{\text{ship}}$ the maximum quantity that can be shipped, while ensuring that the queue remains stable ($\rho < 1$), i.e.

$$\bar{Q}_{\text{ship}} = \sup_{q} \left\{ q \in \mathbb{R}^{+} : q < \frac{S}{2d + T + W_{\text{port}}(q, S, d, K, T)} \right\}$$

since the arrival rate in this queue is determined by q and the average service rate by (2). Unlike the Port Queue, in this case we obtain an implicit definition of $\bar{Q}_{\text{ship}}$ because of the endogenous determination of the queue's service rate. Intuitively, if there are more exporters arriving at port, the port can become congested, so that the wait time W_{port} increases and, as ships are tied up in queue at the port, their effective supply declines.

Then, $\bar{Q}_{\text{ship}}$ satisfies

$$\bar{Q}_{\text{ship}} = \frac{S}{2d + T + W_{\text{port}}\left(\bar{Q}_{\text{ship}}, S, d, K, T\right)} \tag{3}$$

We provide a formal proof in Appendix C.

Intuitively, it takes a single ship $2d + T + W_{\text{port}}\left(\bar{Q}_{\text{ship}}, S, d, K, T\right)$ on average to complete a single round-trip, when the Ship Queue is at full capacity and ships never have to wait for a cargo. Therefore in a unit of time, a ship delivers $1/\left(2d + T + W_{\text{port}}\left(\bar{Q}_{\text{ship}}, S, d, K, T\right)\right)$ units of the good on average and S ships deliver $S/\left(2d + T + W_{\text{port}}\left(\bar{Q}_{\text{ship}}, S, d, K, T\right)\right)$ units of the good.

Proposition 1 summarizes the above analysis:

Proposition 1. *The sequence of queues has maximum capacity Q^{max} which is given by*

$$Q^{max} = \min\left\{\bar{Q}_{ship}, \bar{Q}_{port}\right\} \tag{4}$$

where $\bar{Q}_{ship}$ solves

$$\bar{Q}_{ship} = \frac{S}{2d + T + W_{port}\left(\bar{Q}_{ship}, S, d, K, T\right)}$$

and

$$\bar{Q}_{port} = \frac{K}{T}$$

An obvious implication is that

$$Q(q, S, d, K, T) \leq Q^{max}$$

so that the sequence of queues can deliver at most Q^{max} per unit of time.

Proof. See Appendix C. □

Equation (4) illustrates how the network of three sequential queues contains two constraints or bottlenecks that limit the quantity that can be traded. The first is driven by ship capacity and the second by port capacity. They are both intuitive. Importantly, $Q^{\max}$ is easily computable from basic data sources: the port service time, T, port handling capacity, K, the number of ships S, and the distance traveled d.

Good Demand and Supply

We conclude the description of the environment by specifying demand and supply for the good to be traded.

Consumers' demand in the importing country is given by

$$D\left(p, Y\right)$$

where Y is a demand shifter with $\frac{\partial D(p,Y)}{\partial Y} > 0$ and p is the import price of the good, given by

$$p = p_x + p_T$$

where p_x is the export price and p_T is the transportation price paid to ships.

Producers in the exporting country supply the good according to:

$$q\left(p_x, A\right) \tag{5}$$

where A is a supply shifter with $\frac{\partial q(p_x,A)}{\partial A} > 0$.[14]

Disruptions

Before proceeding to analyze the equilibrium of the model, it is worth discussing how disruptions map in this model. Disruptions occur because of shocks to any of the model's transportation parameters (S, d, K, T) and/or changes in the demand/supply shifters (Y, A).

[14]The exporter supply curve captures not only the cost of production, but also any inland transportation to the exporter's port, as well as packaging costs etc, but not the shipping price, p_T. As a result, one should think of the export price, p_x, as the price at the exporter's port, i.e. Free On Board (FOB).

On the transportation side, disruptions in T capture personnel restrictions, such those experienced during the COVID-19 pandemic, as well as other disruptions to service times such as labor shortages, strikes, crane break-downs, etc.

Disruptions to distance traveled, $1/d$, capture shocks that force ships to travel longer distances, such as the Red Sea crisis, the blockage of the Suez Canal, the recent drought experienced at the Panama Canal, or extreme weather events, such as cyclones and hurricanes that force ships to reroute.

Disruptions to port infrastructure, K capture events that shut down parts of the port, such as extreme weather events, the collapse of the Baltimore bridge, or the closure of the Straits of Hormuz, which effectively shut down access to ports in the Gulf.

Finally disruptions in the number of ships, S, would include the recent US proposal to impose fees on Chinese-built vessels, or the Jones Act that requires US trade to occur via US ships.

At the same time, there may be changes in the demand (Y) or supply (A) of the good, such as the surge in demand during the COVID-19 pandemic, supply disruptions due to the war in Ukraine, COVID-19 related restrictions in mining, or the commodity boom of the 2000s.

Finally, several situations of interest are a combination of different factors; for instance, as we will see, the Great Supply Chain disruption was a combination of demand, supply and transportation (primarily T) shocks.

3.2 Equilibrium

The steady state equilibrium of the above economy is characterized by the following conditions:

$$q\left(p_x, A\right) = q \tag{6}$$

$$D\left(p_x + p_T, Y\right) = Q\left(q, S, d, K, T\right) \tag{7}$$

$$Q\left(q, S, d, K, T\right) = q \tag{8}$$

$$p_T \begin{cases} = c\left(d\right), & \text{if } Q\left(q, S, d, K, T\right) < Q^{\max} \\ \geq c\left(d\right), & \text{if } Q\left(q, S, d, K, T\right) = Q^{\max} \end{cases} \tag{9}$$

Condition (6) states that the arrival rate of exporters, q is determined by the supply curve, (5). Similarly, condition (7) states that the demand for the good equals the quantity delivered. Condition (8)

ensures that quantity arriving at the exporting port equals the quantity delivered at the importing port.

Finally, condition (9) determines the equilibrium shipping price: when the transportation constraint does not bind ($Q(q, S, d, K, T) < Q^{\max}$), then the price is simply equal to the marginal cost of a trip, $c(d)$. When the constraint binds then the price is higher as we show below. We discuss each case in turn.

Case 1: Unconstrained Equilibrium

In the unconstrained case, the equilibrium quantity, $Q(q, S, d, K, T)$ is strictly less than $Q^{\max}$. In this case competition among ships drives the shipping price, p_T, down to the marginal cost of a trip, $c(d)$. We can now solve for the export price, p_x, by equating demand and supply

$$D(p, Y) = q(p_x, A)$$

and setting $p_T = c(d)$, so that, $p = p_x + c(d)$.

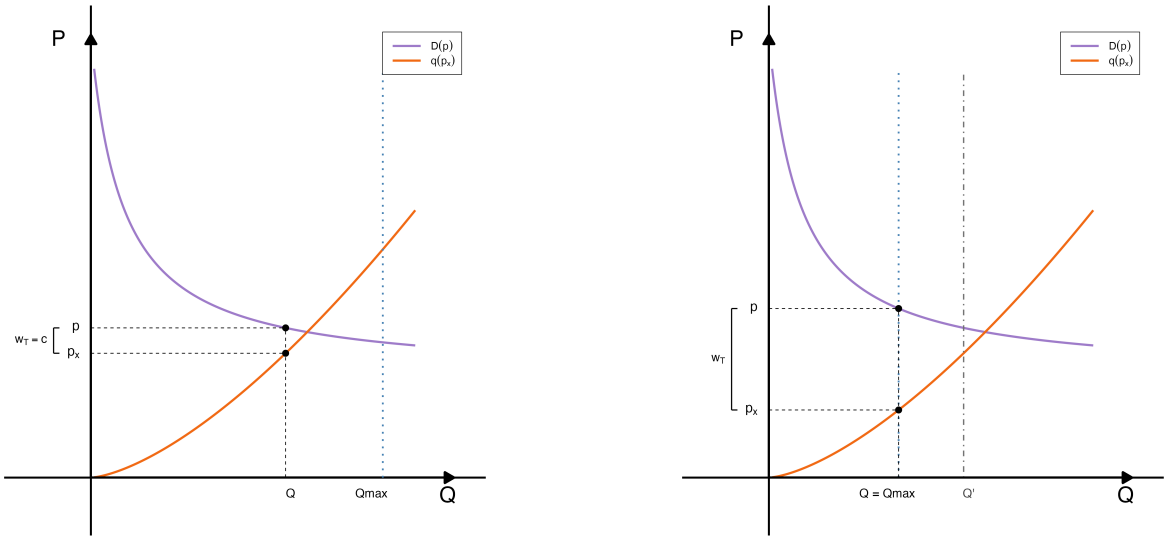


Figure 6: Equilibrium. In the unconstrained case (left panel), the shipping price is equal to the marginal cost of shipping and corresponds to the difference between the price the exporter receives and the importer pays. In the constrained case (right panel) the unconstrained quantity (denoted Q' in the figure) is higher than $Q^{\max}$. The equilibrium quantity is equal to $Q^{\max}$, the exporter price is determined from the intersection of $Q^{\max}$ and the supply curve, while the importer price is determined by the intersection of $Q^{\max}$ and the demand curve. The transportation price is the difference between the two.

As shown graphically in the left panel of Figure 6, the above equation also pins down the equilibrium quantity.

Lemma 1. *In the unconstrained case, the transportation parameters* (S, K, T) *do not affect prices and*

trade, so that for $z \in \{S, K, T\}$,[15]

$$\frac{dp}{dz} = \frac{dp_x}{dz} = \frac{dQ}{dz} = 0$$

Case 2: Port at Capacity

Let us now consider the case where the equilibrium quantity, $Q(q, S, d, K, T)$ equals $Q^{\max}$ and $Q^{\max}$ is determined by the port constraint. In that case, as shown in the right panel of Figure 6, if the transportation price were equal to the marginal cost, $c(d)$, then the equilibrium quantity would be greater than the capacity of the system, $Q^{\max}$. Now transportation prices have to increase in order for the market to clear. In this setup, it is as if ships act as arbitrageurs, earning the difference between the exporter price (net of transport costs) and the importer price.

In particular, since the port constraint binds, the importer price, p, is given by

$$D(p, Y) = \frac{K}{T}$$

where, as discussed above, K/T is the port effective capacity. It is worth noting that in this case, the port parameters, K and T, affect prices (and quantities), but parameters related to the ships, i.e. S and d, do not.

Similarly, we can solve for the export price, p_x, by setting $q(p_x, A) = \frac{K}{T}$ and therefore back out p_T as the difference between the two prices, as shown in the right panel of Figure 6.[16] In other words, the equilibrium quantity is equal to $Q^{\max}$, the exporter price is determined from the intersection of $Q^{\max}$ and the supply curve, while the importer price is determined by the intersection of $Q^{\max}$ and the demand curve. The transportation price is the difference between the two.

Lemma 2. *When the port is at capacity, only the port parameters* $\{K, T\}$ *affect prices and trade, so that,*[17]

$$\frac{dp}{dK} < 0, \frac{dp_x}{dK} > 0, \frac{dQ}{dK} > 0$$

[15] Since the constraint is slack at the baseline, the equilibrium remains unconstrained for marginal changes in $z \in \{S, K, T\}$. Moreover, it is worth noting that d still affects prices through the marginal cost of shipping.

[16] Prior work (Brancaccio et al., 2020, 2023) has modeled shipping price formation via decentralized bargaining under random bilateral meetings. Here, for simplicity, we abstract away from this. Formally, we assume that a clearinghouse trades enforceable, anonymous forward claims to "one unit delivered at the port," with net settlement through the clearinghouse (delivery is verifiable and contracts cannot be renegotiated). This implements the competitive (Walrasian) allocation of the economy with the constraint $Q(q, S, d, K, T) \leq Q^{\max}$, so the equilibrium price is given by $q(p_x, A) = D(c(d) + p_x, Y)$ when the constraint is slack, and by $q(p_x, A) = Q^{\max}$ when it binds.

[17] As above, a marginal change does not affect whether and which constraint binds, as long as the unconstrained equilibrium quantity is strictly above $Q^{\max}$.

$$\frac{dp}{dT} > 0, \frac{dp_x}{dT} < 0, \frac{dQ}{dT} < 0$$

while for $z \in \{S, d\}$,

$$\frac{dp}{dz} = \frac{dp_x}{dz} = \frac{dQ}{dz} = 0$$

Case 3: Ships at Capacity

When the ship constraint binds, the import price, p, is given by

$$D(p, Y) = \bar{Q}_{\text{ship}}$$

where $\bar{Q}_{\text{ship}}$ is now given by the solution to (3) and we can similarly solve for the export price, p_x, by setting $q(p_x, A) = \bar{Q}_{\text{ship}}$.

Unlike the case where the port constraint binds however, now marginal changes to *any* transportation parameter, S, d, K or T, affect prices and quantities, as they change $\bar{Q}_{\text{ship}}$ given by (3). Intuitively, an increase in the number of ships, S, directly increases ship capacity. Similarly a reduction in expected travel time, d, allows ships to complete more trips per unit of time. Interestingly, however, changes in K and T also affect the available shipping capacity: for instance, a lower service time, T, implies that ships spend less time at the port and as a result this increases the effective supply of shipping by increasing the number of possible trips per unit of time. The same is true if the port can handle more ships at a time so that K expands and ship wait time at the port falls.

Lemma 3. *When ships are at capacity, all transportation parameters* $\{S, d, K, T\}$ *affect prices and trade. For* $z \in \{S, K\}$,

$$\frac{dp}{dz} < 0, \frac{dp_x}{dz} > 0, \frac{dQ}{dz} > 0$$

while for $z \in \{d, T\}$,

$$\frac{dp}{dz} > 0, \frac{dp_x}{dz} < 0, \frac{dQ}{dz} < 0$$

Proof. See Appendix C. □

Model Implications This model has sharp implications for trade and prices. It predicts that there are two underlying regimes; one where the economy is unconstrained and prices and quantities are determined by demand and supply, and one where the transportation sector is at capacity and determines quantities. When a transportation constraint binds, shipping prices spike. As it turns out, this is a prevalent feature

of the data.

The top panel of Figure 7 plots the Baltic Exchange Dry Index (BDI), which is a commonly-used index of dry bulk shipping prices. As shown in the figure, prices are largely flat for long periods of time, such as for most of the 2010s, but there are periods where they rise sharply, stay high for a few years, but then return to their previous levels. For instance, prices increased substantially in Great Supply Chain Disruption, approximately tripling from 2021 to 2022 and then slowly returning back to their pre COVID-19 levels. A larger and more prolonged increase in dry bulk prices took place between the mid 2000s, until approximately the end of that decade, coinciding with the Commodity Supercycle, driven mostly by China's demand for raw materials. Consistent with this, interviews with shipowners suggest that shipping profits are concentrated in these relatively short periods where shipping prices exceed their long-run level, with minimal earnings during the extended intervals between such episodes (Kalouptsidi, 2014). Our framework captures this cyclical profit structure.

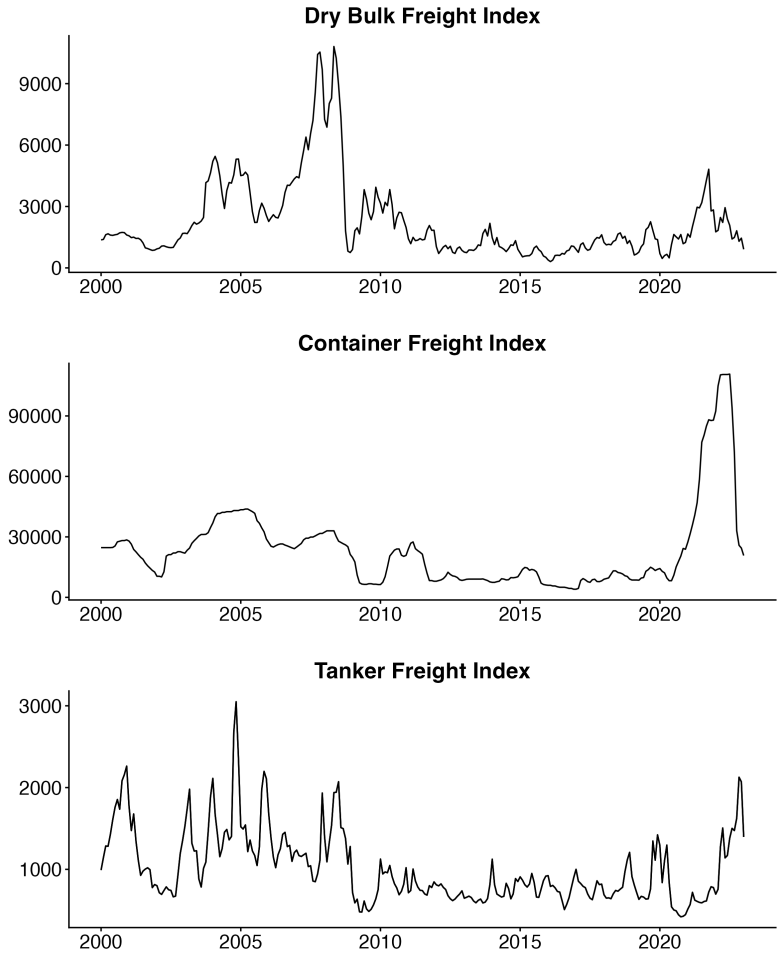


Figure 7: Shipping Prices. Top panel: Baltic Exchange Dry Index (BDI). Middle: Containerized Freight Index. Bottom: Tanker Freight Index.

This pattern does not seem to be unique to dry bulk shipping prices. The middle and bottom panels of Figure 7 plot price indices for containership and oil tanker freight rates respectively, the other two major segments of oceanic shipping. While the timing and magnitude of the price increases differ across the three segments, the pattern is the same, with long periods of largely flat prices, punctured by periods of sharp increases, where prices often more than triple. It is worth noting that the increase in prices cannot be explained by increases in operating costs, e.g. by bunker (fuel) price increases. This can be seen from the patterns in Figure 7: even though bunker (fuel) prices for all three ship categories co-move over time, the corresponding shipping prices peak during different time periods and with different intensities.

The equilibrium characterized above highlights the role of the transportation sector capacity, which gives policymakers a transparent, high-frequency tool for anticipating and monitoring these shipping cycles. The transportation sector's capacity in (4) can be computed from widely available data, providing an immediate indicator of the sector's resilience.

The shipping price patterns in Figure 7 are exactly as predicted by our model. In Section 4 we calibrate this model and check whether the data is overall consistent with the presence of two regimes (constrained vs. unconstrained).

Discussion We close this section with a brief discussion on some of our assumptions and related caveats.

Our model is one of two countries and uni-directional trade. The latter is justified by the nature of bulk trade, as trade in commodities is largely imbalanced, with a small number of exporting countries specializing in one commodity and a large number of importers. The assumption of two countries, which is common in the literature, is made to simplify the exposition and clearly illustrate the economics of the two regimes (constrained and unconstrained) and provide a simple and easy-to-compute measure of the transportation sector's capacity. For instance, it can be interpreted as a Home versus Rest-of-the-World representation, which is often the relevant abstraction from the perspective of a domestic policymaker. In addition, as discussed above, trade in many commodities is highly concentrated among a small number of exporting countries, making it difficult for importers to reroute purchases following a supply disruption. We thus expect the key mechanisms to go through with multiple countries; however the analysis quickly becomes significantly more complex and is left for future research.[18]

In our setup, the importing port is unconstrained; this is done for expositional simplicity. It is straightforward to similarly model the importing port. This would essentially add one more (analogous)

[18]To give one example, with multiple countries we need to embed a model for how ships choose where to sail after unloading at some location (such as the one in Brancaccio et al., 2020). This significantly complicates the analysis.

port bottleneck without changing the economic intuition.

In addition, our setup assumes that ships but not ports can reap the rents when the system binds. This is motivated by the observation that, unlike shipping prices, port prices are remarkably stable (see Brancaccio et al., 2024) due to the semi-public ownership structure of ports. It turns out this has important policy implications which we discuss in the Conclusion.

Finally, it is a straightforward extension to allow supply (or demand) to depend on total time in the transportation system, in addition to price, thereby capturing the value of time related to inventory costs and other factors. In fact, in our empirical exercise below we have calibrated this version of the model using estimates for value of time from Brancaccio et al. [2024]. We found this to have a very small impact on the results.

4 Empirical Implementation

We bring the model to the data to understand the role of the transportation sector. Of course, this is complicated by other factors, such as contemporaneous demand and supply shocks, as well as fuel costs. Our empirical strategy is designed to differentiate all these sources of variation and distinguish the role and impact of each relevant factor on prices and trade and more importantly, how they interact.

Given the heavily imbalanced nature of trade in raw materials, we divide the world into countries that import raw materials and countries that export them. Our model accordingly features two countries, one standing in for importers and one for exporters.

In this section we discuss how we calibrate the model. We then illustrate the economics at play by considering the impact of a port service time disruption. In Section 5 we turn to our main question of measuring the capacity of the transportation system and exploring whether and when it became binding during the last twenty years, affecting prices and trade.

4.1 Calibration

We calibrate the model parameters using our data and some estimates from the literature. We first calibrate the model for the most recent period, 2013 to 2023, and then with some additional assumptions (discussed in Appendix D separately), we calibrate it to the 2003 to 2009 period. Our model consists of two blocks: one that governs the transportation sector and determines $Q^{\max}$, and another that governs demand and supply for traded raw materials.

Let us begin with the transportation parameters (S, d, K, T). A period is a quarter. The fleet size S is observed directly in the Clarksons dataset biannually. We compute the average port service time, T, quarterly, using the ship tracking (AIS) data. This dataset is also used to set the average travel duration, d, to 2.9 weeks, which is the average travel time.

In order to calibrate the port handling capacity K, we follow Brancaccio et al. [2024] and use the AIS data for every exporting port to compute the maximum number of ships that a port can simultaneously handle, while there are ships waiting in anchorage. We condition on ships waiting, to ensure that there is no unused capacity inside the port. We use the 75th percentile of this measure to define the port-specific K, and then aggregate over all ports. To account for improvements in port infrastructure, we construct separate measures for the periods 2013-2016, 2017-2019, and 2020-2023.

Next, we estimate the marginal cost of a trip every quarter as the sum of bunkering (i.e. fuel) prices, average port charges, and a measure of fixed operating costs. The first component is simply equal to the observed fuel cost times the average distance d. The second component, port charges, is set equal to 3$ per ton (Brancaccio et al., 2024). For the final component of operating costs, we exploit the fact that under time-charter contracts, shipowners bear the vessel's fixed costs—including crew, maintenance, and repairs. We calibrate these operating costs to the lowest level of time-charter prices based on Clarksons data; see Appendix D.

Finally, we turn to the demand and supply of commodities. We assume the demand curve is given by

$$\frac{\sigma - 1}{\sigma} \frac{Y_t}{p} \tag{10}$$

where $\sigma > 1$, [19] while the exporter supply curve is given by

$$A_t p_x^{\beta}$$

where $\beta > 0$. We allow both the demand and supply shock A_t and Y_t to vary quarterly.

We calibrate α and β from the literature. Namely, we set the supply elasticity, β, equal to 2, following recent estimates by Garcia-Lembergman et al. [2025] who focus on fossil fuels (coal is one of the three major commodities transported by dry bulk ships), as well as results in Broda et al. [2008] who provide supply elasticities for various commodities, including iron ore and grain (the other two major commodities). Similarly, we set the demand parameter, σ, equal to 4, following Simonovska and Waugh [2014].

[19] Demand (10) can be microfounded as the demand for an imported input used by monopolistically competitive firms.

To recover the demand and supply shocks A_t and Y_t we leverage data on quantities traded Q_t, as well as transportation and exporter prices for raw materials, $p_{T,t}$ and $p_{x,t}$, and compute

$$Y_t = \frac{\sigma}{\sigma - 1} Q_t \left(p_{x,t} + p_{T,t} \right)$$

and

$$A_t = \frac{Q_t}{p_{x,t}^{\beta}}$$

We normalize prices, quantities, port handling capacity, and fleet size by shipment size, which is defined as that of a median vessel (a Handymax of 57,000 DWT). We also construct quarterly port service times for the median vessel by residualizing raw service times with respect to ship size, thereby accounting for systematic variation in service times across vessel sizes. All details in Appendix D.

Table 2 summarizes the parameters and the moments we use to calibrate them.

Parameter	Target
Total Available Ships, S_t	Biannually, Number of dry bulk ships
Cost of Trip, c_t	Quarterly, Fuel cost + fixed operation cost + port dues
Trip Duration Parameter, d	Average trip duration (2.9 weeks)
Port Service Time, T_t	Quarterly, Average port service time
Port Handling Capacity, K_t	Tri-annually, Max number of ships handled simultaneously when congested
Export Supply Curve Parameter, β	$\beta = 2$, Garcia-Lembergman et al. (2025) and Broda et al. (2008)
Demand Parameter, σ	$\sigma = 4$, Simonovska and Waugh (2014)
Export Supply Shock, A_t	Quarterly, Producer price+quantity
Demand Shock, Y_t	Quarterly, Commodity price+quantity

Table 2: Calibration of Model Parameters.

4.2 Model Illustration: Port Service Time Disruptions

We illustrate the model by considering how port service time disruptions affect prices and quantities. We focus on T disruptions since they played a crucial role during the Great Supply Chain Disruption, as we show later.

Using our estimated parameters for the second quarter of 2019, the three panels of Figure 8 show how

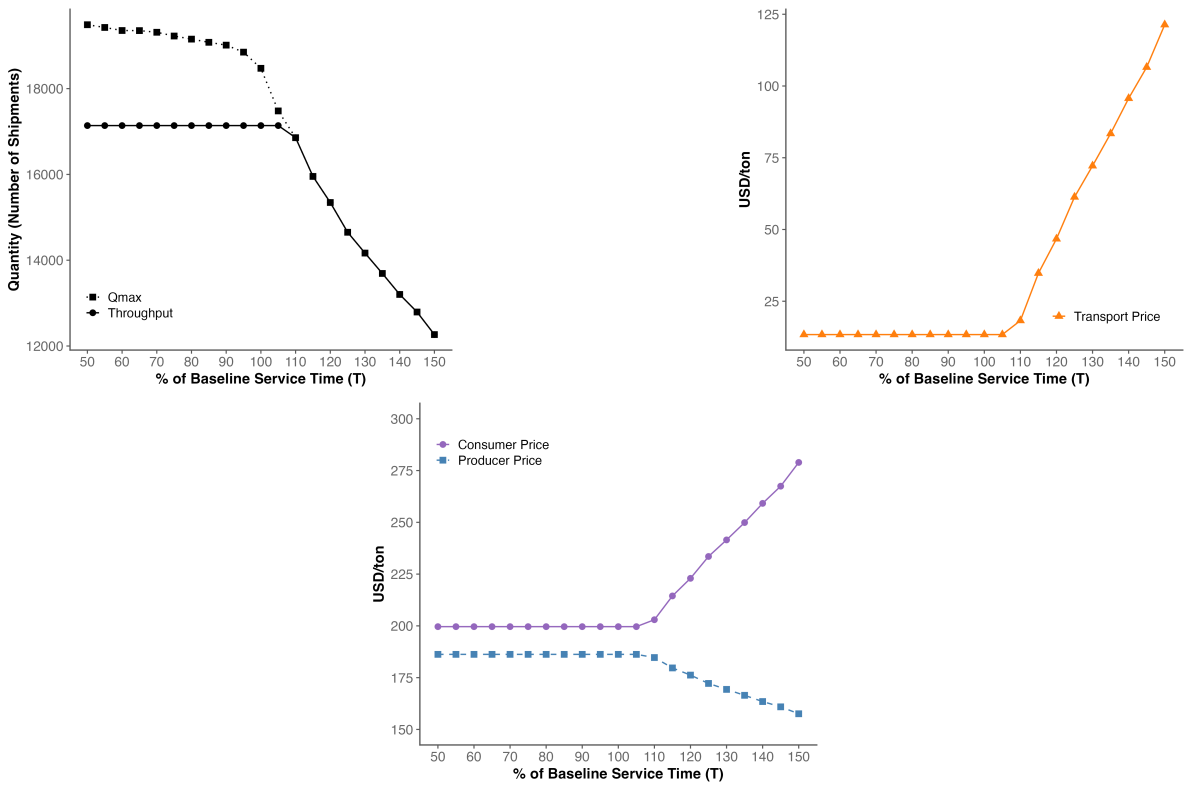


Figure 8: Impact of service time disruptions: simulation using estimated model parameters, where the baseline (100%) is at the second quarter of 2019.

changes in service time, T, affect equilibrium quantity, Q, and transportation capacity, $Q^{\max}$ (top left panel), shipping price, p_T (top right panel) and import and export prices, p and p_x respectively (bottom panel).

First off, it is important to note that at the Q2-2019 baseline, before the Great Supply Chain Disruption, the transport sector was operating with slack ($Q < Q^{\max}$) and transportation capacity $Q^{\max}$ was determined by port bottlenecks ($Q^{\max} = K/T$). As shown in the top left panel, a reduction in service time would further relax the port constraint and raise $Q^{\max}$. However, neither equilibrium quantity, nor any of the prices would change. Indeed, service time T affects the transport sector capacity $Q^{\max}$ but, when the transport sector is operating below capacity, Q^{max} does not directly affect prices and quantities (Lemma 1). Moreover, if T fell below approximately 80% of its baseline value, the port capacity would increase so much that the ship constraint would instead determine $Q^{\max}$. Further reductions in T would still increase $Q^{\max}$, but the rate of increase would be lower, since T matters less for the ship constraint than it does for the port constraint.

The opposite case—an increase in service time above baseline, as occurred during in 2021-2022 and shown in Figure 15 in Appendix A—is more consequential. Here $Q^{\max}$ falls sharply: an increase in service

time T of just over 10% is enough to make the port constraint bind. Once it binds, further increases in T reduce trade, as Q is now pinned to the declining capacity $Q^{\max}$. The second panel shows that this coincides with a marked change in shipping prices p_T, which begin to rise. The third panel traces how this feeds through to prices: import prices p increase, while export prices p_x fall.

5 Transportation Bottlenecks in the Past Twenty Years

As discussed above, our framework provides a model-based estimate for the capacity of the global transportation sector. We begin by measuring this global transport capacity, $Q^{\max}$ between 2004 and 2023. Then, we explore how the state of the sector mediated the impact of demand and supply shocks on the economy during this period. We split the analysis in two, the years around the Great Supply Chain Disruption (2013-2023) and the Commodity Supercycle of the 2000's (2004-2009).

The Great Supply Chain Disruption, 2013-2023

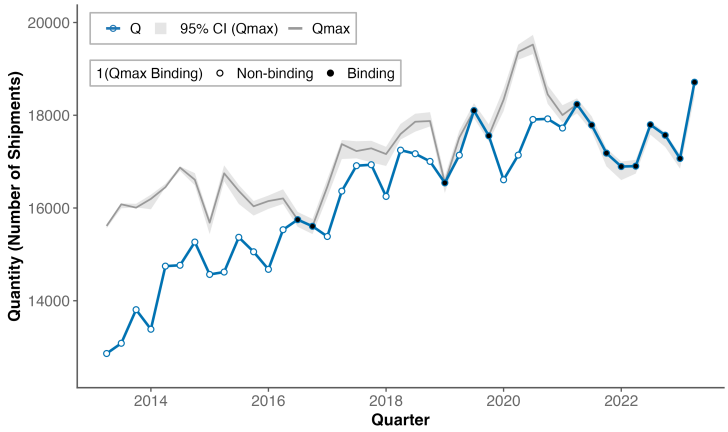


Figure 9: Transportation capacity, $Q^{\max}$, and quantity traded, Q, 2013-2023. Confidence bands are computed based on 100 bootstrap samples, as described in Appendix D.

Figure 9 shows our estimate for the transport sector capacity $Q^{\max}$, as well as the observed quantity shipped, Q. $Q^{\max}$ is obtained from data on port service time and port handling capacity, as well as on the effective supply of ships, as described in Proposition 1. In the beginning of the period, the ship constraint determined $Q^{\max}$; after 2014, the port constraint became the relevant one.[20] At the start of this period, around 2013-2014, the dry bulk shipping sector is able to transport up to around 17,000 ship-sized shipments (973 million tons) every quarter and this slowly trends upward, due to incremental

[20]This is consistent with the results in Dunn and Leibovici [2025], who find that ship investment and ship utilization decisions in container shipping are not a leading driver of price dynamics during the COVID-19 pandemic.

improvements in port service times, as well as expansions in in port infrastructure. However, starting around 2021, we see a substantial decline in $Q^{\max}$. Around this time, port closures and COVID-19 related personnel restrictions at ports led to sizable increases in service time, as shown in Figure 15 in Appendix A. This increase in T, led to a reduction in port handling capacity and the decline in $Q^{\max}$ shown in Figure 9.

Quantities transported increase steadily throughout this period. Until 2021, quantity shipped is mostly below the transportation sector's capacity ($Q < Q^{\max}$). Beginning in 2021 however, $Q^{\max}$ becomes binding and this restricts trade in commodities for several quarters. Strikingly, as can be seen in the right panel of Figure 20 in Appendix A, shipping prices spike when the transportation sector has no slack so that throughput reaches $Q^{\max}$; this acts as a model validation, since we do not use prices when constructing the system's capacity $Q^{\max}$.

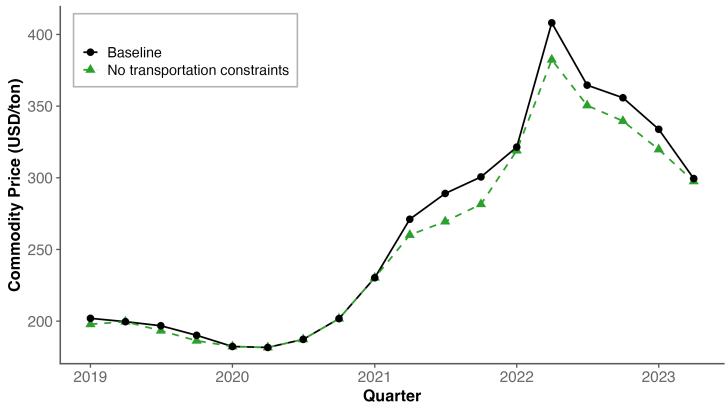


Figure 10: Commodity prices: baseline and counterfactual with no transportation constraints 2019-2023.

What was the impact of the transportation bottlenecks on prices and trade during the Great Supply Chain Disruption? To answer, we use the estimated demand and supply curves and simulate a counterfactual world where the transportation sector's capacity is higher, so that $Q^{\max}$ no longer binds. Figure 10 shows the resulting commodity prices and compares them with the baseline ones, while Figure 17 in Appendix A shows the resulting quantities traded, as well as the baseline ones. As shown in Figure 10, commodity price inflation would have been three percentage points lower in the absence of transportation bottlenecks during 2021 to 2023, relative to 2019, which corresponds to 13% of the total increase during that time.

Digging deeper into the drivers of equilibrium prices, we find that, first, as shown in Figure 16 of Appendix A much of the increase is, not surprisingly, due to a large increase in demand for commodities, Y. This is consistent with much of public discourse during those times, where demand for consumer goods

spiked, naturally also driving demand for inputs. More interestingly however, starting in 2021 there is also a considerable reduction in supply, A, driven by COVID-19 related lockdowns in many exporting countries, mine closures, as well as the war in Ukraine that greatly affected grain, fertilizer and coal exports.

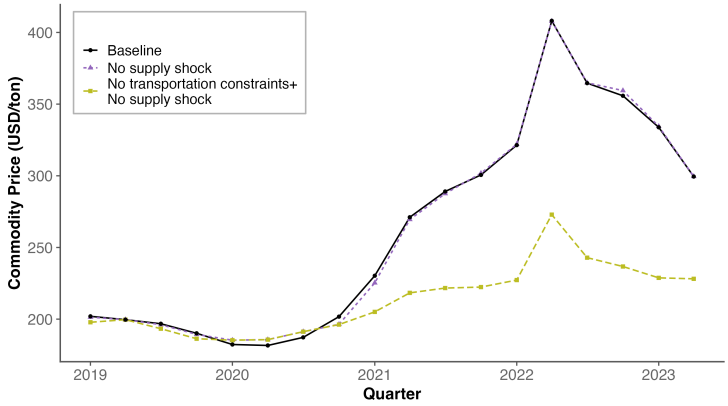


Figure 11: Commodity prices: "No supply shock", we set the supply shifter, A, between 2021-2023 to its average 2019 level; in "No transportation constrains + No supply shock", we set the supply shifter, A, to its average 2019 level, and remove transportation constraints.

How did these three factors (demand shock, supply shock, transportation bottlenecks) interact to produce the rise in prices? Our framework can provide an answer.

Figure 11 shows commodity prices if we shut down the negative supply shock, by setting the supply shifter, A, to its average 2019 value (similarly Figure 18 in Appendix A shows commodities traded if there were no supply shock). Somewhat surprisingly, when we shut down the negative supply shock, there is very little change for both commodity prices and quantities traded. Why? Without slack in the transportation sector, the higher supply/production cannot be transported to consumers, preventing trade from increasing and leading to similar price increases. Indeed if we simultaneously increase the transportation sector's capacity and shut down the supply shock at the same time, there is a very sizable reduction (72%) in the increase in commodity prices, while commodity trade increases by almost 50%.

This suggests that during the Great Supply Chain Disruption price increases seem to have been inevitable, as either the negative supply shock or the binding port bottlenecks would have caused them.

The result carries a sharp policy implication. While the transportation constraint binds, measures that boost foreign commodity supply — subsidizing production, releasing stocks — do little to ease upward pressure on prices, because they relax a margin that is not binding. This may explain in part why big importers often invest in foreign transportation infrastructure, as China has done extensively in Africa

and Southeast Asia.

The Commodity Supercycle, 2004-2009

We now consider another period when the transportation sector came into the spotlight. As shown in Figure 7, in the mid to late 2000s shipping prices for commodities exploded. This surge was driven largely by China's rapid industrialization and urbanization, which led to a substantial increase in demand for raw materials such as iron ore, coal, steel and metals needed to build factories, transportation networks, housing, etc. We now use our framework to illustrate that at the same time, the supply of shipping capacity struggled to keep pace with this growth in trade, playing a key role in the sharp rise in prices.

There is no dependable AIS data for this period.[21] However, as we discuss in Appendix D, using somewhat coarser data and some conservative assumptions, we are able to calibrate our model for this earlier period.

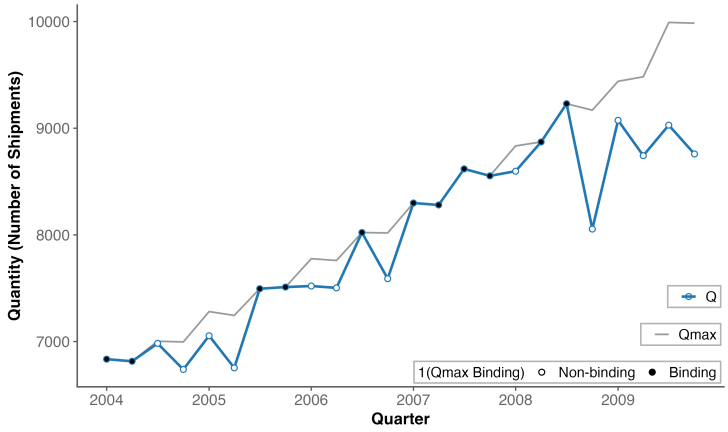


Figure 12: Transportation capacity $Q^{\max}$, and quantity traded, Q, 2004-2009.

Figures 12 shows our estimate for $Q^{\max}$, as well as the observed quantity shipped.Unlike the later bout of price spikes discussed above, it is not the port constraint that determines $Q^{\max}$, but rather the ship constraint, which, as shown in the figure, was often binding during this period. During this period there was a massive increase in orders for new ships, which, because of time-to-build in shipbuilding, did not come online until the early 2010s and would continue to enter the fleet until the mid-2010s, well after the commodity boom had ended and demand for shipping services had largely ebbed following the Great Recession of 2009 (Kalouptsidi, 2014). As a result, the dry bulk fleet increased by 140% in the number of ships and by 239% in deadweight capacity in the 2002-2023 period and the ship constraint did not bind

[21] While most dry bulk ships were required to use AIS transponders starting from December 31st, 2004, the necessary coverage to pick up these signals for the entire globe was not in place until almost a decade later.

again after 2013. As predicted by our model, shipping prices increased manyfold before collapsing at the Great Recession, consistent with the periods when the ship constraint was binding, as shown in the left panel of Figure 20 in Appendix A.

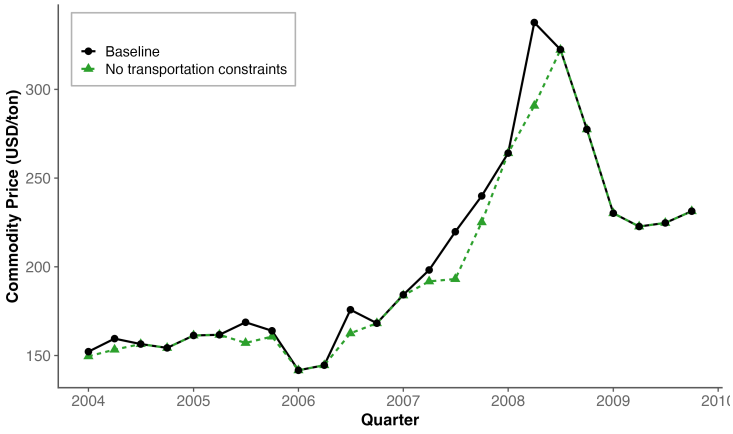


Figure 13: Commodity prices: baseline and counterfactual with no transportation constraints, 2004-2009

Figure 13 shows the increase in commodity prices during this period, as well as counterfactual prices, when we simulate a world with higher transportation capacity, so that the ship constraint would not bind (see also Figure 19 in Appendix A for the corresponding quantities traded). The impact of transportation bottlenecks is significant, as they are responsible for up to 19% of the increase in commodity prices during the period 2007 to mid-2008, relative to 2006.

Our results demonstrate that transportation infrastructure can have aggregate consequences, shaping the pattern of international trade and the availability of resources around the world. Policies that affect these margins, such as the semi-public ownership structure of ports, shipbuilding subsidies, or how congestion rents are captured across ships and ports, thus carry first-order implications for global commodity markets. In the Conclusion we discuss the implications of our analysis for investment in transportation.

Impact of Other Disruptions: The Red Sea Crisis

We can also use our setup to predict the impact of other shocks and disruptions. As is evident from the previous discussion, the effect of any shock on the economy depends on how close the transportation sector is to hitting its capacity, $Q^{\max}$, but also on whether $Q^{\max}$ is determined by port vs. ship bottlenecks. For instance, a shock to the number of ships may matter little if the port constraint is the one determining $Q^{\max}$.

To see this, we examine the impact of the Red Sea crisis, which has received significant attention in

the press.[22] The Red Sea crisis involves attacks on cargo ships traveling through one of the world's most important sea routes. As a result, a number of dry bulk vessels have chosen to avoid the Red Sea and travel around Africa (Cape of Good Hope) instead of the Suez Canal. The attacks on cargo ships began in October 2023, a few months after our data sample ends. We are able however to use our framework to both predict, but also understand, the impact of this crisis on commodity prices and trade.

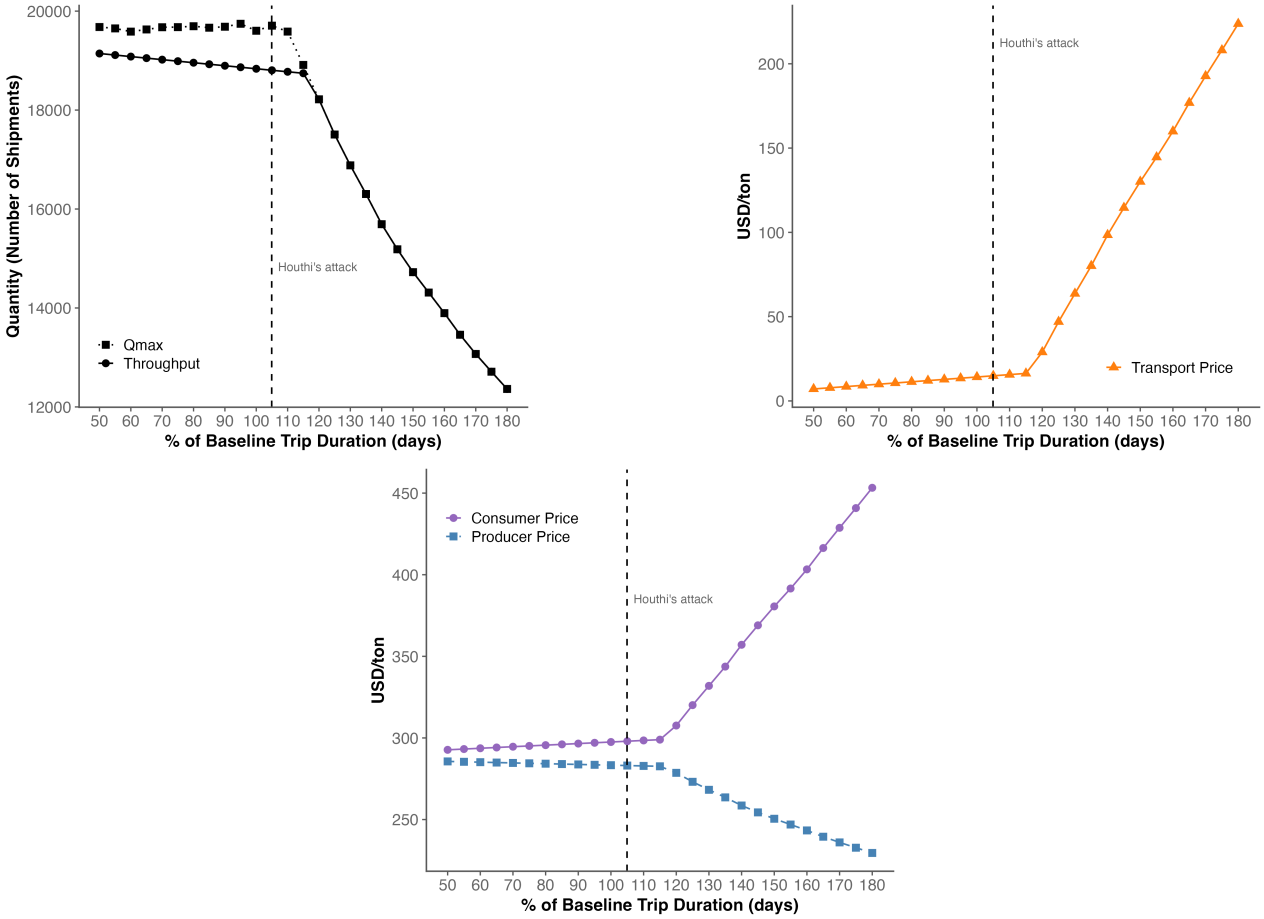


Figure 14: Impact of disruptions in trip duration because of the Red Sea Crisis: simulation using estimated model parameters, where the baseline (100%) is at the second quarter of 2023. The vertical dashed line represents the increase in average travel time due to ship rerouting. We calibrate counterfactual trip duration during the Red Sea crisis as $90\% * 2.9$ weeks $+ 10\% \times (2.9$ weeks $+ 10$ days), where only the Asia/Oceania–Europe component is assumed to experience the additional delay.

More specifically, we use our estimated model for the last period in our data sample, as well as data on the corresponding increase in ships' travel time (d) due to the crisis, to simulate its impact. Figure 14 shows how capacity $Q^{\max}$, quantity traded, and prices change as a function of ships' travel time. The baseline (100%) represents Q2-2023, while the vertical dashed line represents the increase in average travel time due to the attacks.

[22]For example, see The Economist [2023] and Jones and Wright [2024].

It is important to note that, as discussed above, port bottlenecks—rather than ship bottlenecks—determine $Q^{\max}$. As a result, as shown in the top left panel of Figure 14, increases in d that are less than 15% do not affect $Q^{\max}$, since that is determined by the port constraint. Increases larger than 15% do lead to a reduction in $Q^{\max}$, since at that point it is the ship constraint that determines $Q^{\max}$, but nonetheless $Q^{\max}$ does not bind until d increases by 20%.

As a result, any increases in travel time less than 20% have a relatively small impact on prices and trade, since the transportation constraint is slack. Trade and prices are still affected because increased travel time raises the fuel costs and therefore the cost of a trip, $c(d)$, but the impact is relatively contained.[23] This is in part because the attacks primarily affected the Asia/Oceania to Europe route, which as opposed to containerships, is not as important in bulk commodities.[24]

This example underscores how the impact of a shock, whether stemming from the transportation sector (i.e. disruptions in T, d, K or S), or shocks to the supply or demand of commodities (shocks to Y and A), depends on the state of the transportation sector. More specifically, it depends both on how close the transportation sector is to capacity, $Q^{\max}$, but also on whether $Q^{\max}$ is determined by the port or the ship constraint.

This model can also be used to conduct similar exercises for other types of shocks, such as the recent closure of the Strait of Hormuz. The blockade incapacitated exporting ports and left many ships effectively trapped in or around the strait. Indeed, the strait's strategic role during the recent tensions illustrates the broader importance of transportation infrastructure. Since tankers operate similarly to bulk carriers, one could use tanker data and the same set of calculations to estimate the effect of disruptions in the Strait of Hormuz on oil prices.

6 Conclusion

This paper develops a tractable framework of the transportation sector to identify its bottlenecks and quantify the effects of disruptions on prices and international trade. It also provides an easy to compute measure of the shipping sector's capacity which policymakers can use to evaluate its resilience. Port bottlenecks increased commodity prices during the Great Supply Chain Disruption, although a contem-

[23] In practice, we might also expect an increase in the cost of insurance for the ships that do travel through the Red Sea, which would be captured as an increase in c_0.

[24] The attacks increased travel duration between Asia/Oceania and Europe by approximately 10–14 days (source: Atlas Institute of International Affairs, accessed here). However this route accounts for less than 10 percent of dry bulk trade. This is consistent with AXSMarine [2024], which shows that although the Red Sea crisis sharply reduced aggregate Suez Canal traffic, the effect was concentrated most strongly in container shipping, and dry bulk was the least affected major vessel segment.

poraneous negative supply shock mitigated these effects, while shortages of shipping capacity played a key role during the 2000s Commodity Supercycle.

Our results demonstrate that transportation infrastructure can have aggregate consequences. Policies that affect these margins, such as the semi-public ownership structure of ports, shipbuilding subsidies, or how congestion rents are captured across ships and ports, thus carry first-order implications for global commodity markets. As an example, a striking feature of the data is that while shipping prices exhibit dramatic increases (that coincide with periods of binding transportation bottlenecks), port prices remain flat over this period (see Brancaccio et al., 2024), even when the relevant bottleneck is ports rather than ships. Viewed through the lens of the previous analysis, this leads us to speculate that although there may be ample incentives for shipowners to invest in the fleet, since they reap profits when freight rates spike, there may be possibly insufficient incentives to invest in port infrastructure.[25] There are thus first-order policy implications for port infrastructure investment (Brancaccio et al., 2024), as well as port pricing, which we explore in ongoing work.

[25]Indeed consistent with this, in a 2026 article, the CEO of a major shipping company stated "that now bottlenecks were emerging in ports." Additionally, he "had repeatedly warned about the dangers of overcapacity in the container industry after rivals bought large numbers of new vessels with their bumper profits following Covid. [...] "Since the financial crisis, we have underinvested in infrastructure. I don't think there is a quick fix today. What it's going to take to catch up on 15 years of under-investment is a bit of a catch-up cycle on infrastructure investments."."

Appendix

A Additional Tables and Figures

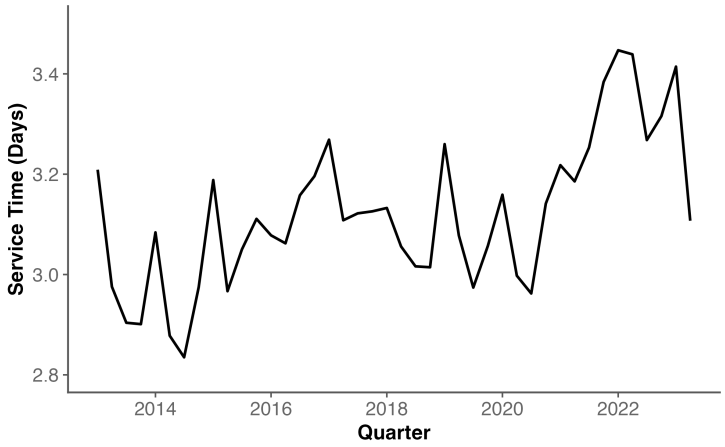


Figure 15: Average port service time for exporting ports, T, 2013-2023.

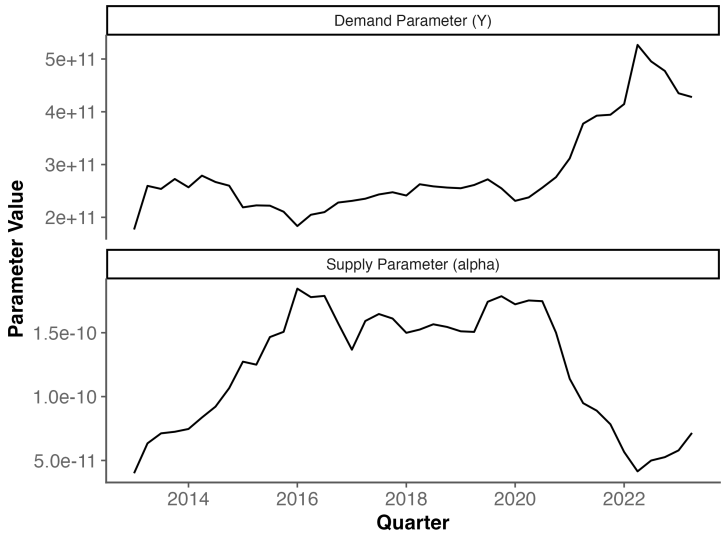


Figure 16: Estimated demand (Y) and supply (A) shifters, 2013-2023, obtained as described in Section 4.

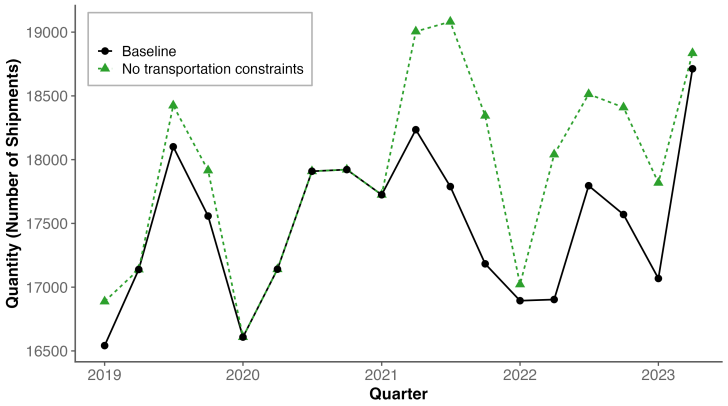


Figure 17: Commodity quantities traded: baseline and counterfactual with no transportation constraints 2019-2023.

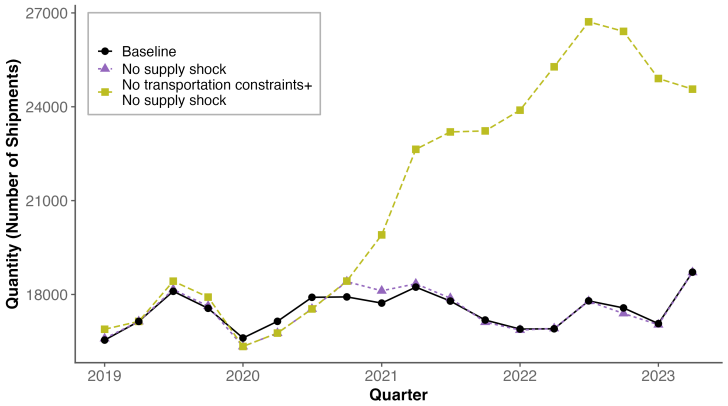


Figure 18: Commodity quantities traded: in "No supply shock", we set the supply shifter, A, between 2021-2023 to its average 2019 level; in "No transportation constrains + No supply shock", we set the supply shifter, A, to its average 2019 level and remove transportation constraints.

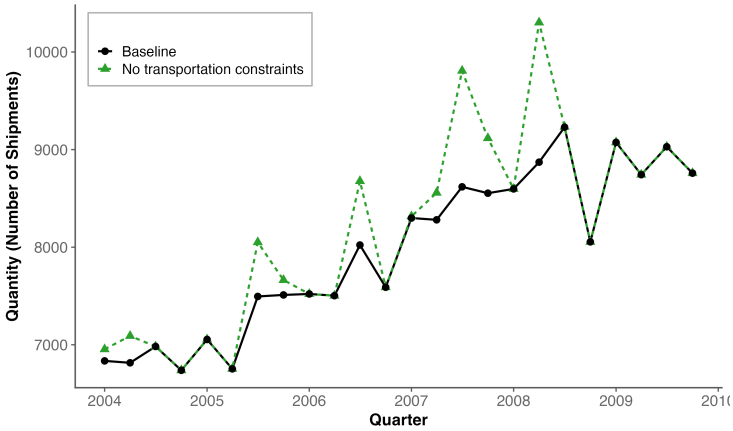
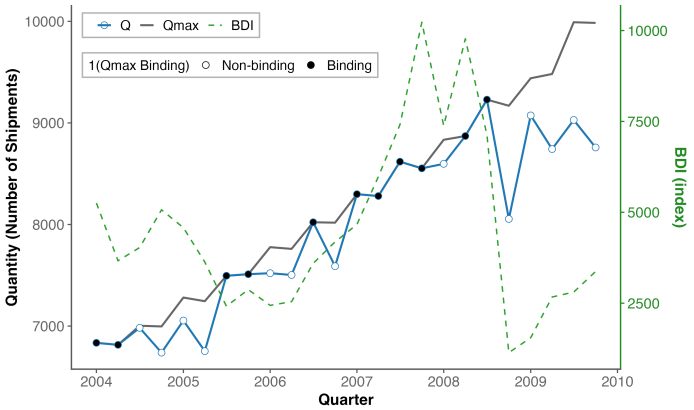


Figure 19: Commodity quantities traded: baseline and simulation with no transportation constraints, 2004-2009.

Panel A: 2004 - 2009



Panel B: 2013 - 2023

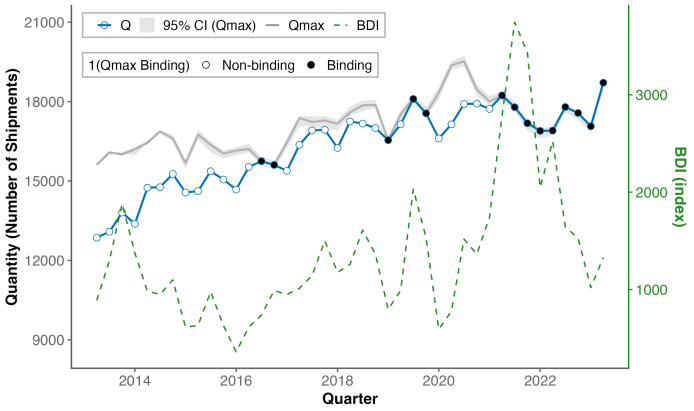


Figure 20: The first panel plots transportation capacity, $Q^{\max}$, the observed quantity traded, Q, as well as the Baltic Dry Index (BDI) between 2013 and 2023. The second panel plots transportation capacity, $Q^{\max}$, the observed quantity traded, Q, as well as the BDI between 2004 and 2010.

B Descriptive Evidence

In this section we discuss the construction of the main variables used in the descriptive evidence in Section 2.3; we then provide details on the implementation and validation of the regressions presented in the same section.

B.1 Construction of Key Variables

We provide details on how the key variables of our analysis are constructed.

Transportation Price, p_T. We construct a measure of shipping prices, p_T, for each exporter - importer - commodity triplet, using granular fixture (or trip contracts) data. We have detailed data on a sample of fixtures, that is prices paid to hire a vessel for a specific trip between specified ports quoted in U.S. dollars per ton (known as voyage charter rates), from Clarksons. However, these data are only available for 2000–2015 and 2020–2023, and we do not observe fixtures for all country pairs. We therefore need to impute p_T for missing years and trading pairs.

To do so, we first assign each country to one of 15 world trading regions, defined following Brancaccio et al., 2020. We then aggregate the fixture data to the exporter region–importer region–month level. Next, we regress observed charter rates on a set of variables capturing key determinants of shipping prices, including the Baltic Dry Index (BDI), fuel costs, voyage distance, average vessel size, exporter-region fixed effects, and importer-region fixed effects. We use the estimated relationship to predict transport prices for missing observations. The regression, reported in Table 21 together with its fit, yields an R^2 of 0.89, indicating that these variables have strong predictive power for p_T.

Export Price, p_x. Our primary data source for export prices is the Harmonized System (HS) bilateral trade data from the United Nations COMTRADE dataset, which is available at a monthly frequency. As this dataset reports both quantities and values of bilateral trade, we use the ratio of the two to obtain the bilateral export price, or unit value.[26] We focus on the following major commodities recorded in the AXS data: coal, grains (wheat, soybean, corn), iron ores, fertilizers, and cement.[27]

[26] Following Amiti et al., 2024, we use exporter-reported value on a free-on-board basis (FOB) to measure prices, rather than importer-reported value, because certain importers do not consistently report trade values of imports on a cost, insurance & freight basis (CIF). We also drop observations with unit values below the 1st percentile or above the 99th percentile within each commodity.

[27] The preferred measure of prices is based on 6-digit level COMTRADE data. In particular, we use 6-digit level data for coal (270112, 270119), grains (100590 for corn, 100199 for wheat, 120190 for soybean), ores, and cement. For fertilizers, we rely on 4-digit level data (3102, 3104, 3105) because each of the 4-digit code consists of multiple 6-digit level codes, and we are not able to map AXS to each of them separately.

Model:	(1)
Variables	
Voyage days	-0.174***
	(0.059)
Fuel cost per day-tonne	-0.697
	(3.55)
BDI	0.002***
	(0.0002)
Ave. Ship type = Handymax	8.91***
	(1.02)
Ave. Ship type = Handysize	14.4***
	(1.22)
Ave. Ship type = Panamax	3.65***
	(0.899)
Ave. Ship type = VLBC	-1.17**
	(0.493)
Voyage days × Fuel cost per day-ton	1.62***
	(0.166)
Voyage days × BDI	0.0001***
	(1.1×10^{-5})
Fixed-effects	
Exporter Region FE	Yes
Importer Region FE	Yes
Fit statistics	
Observations	4,881
R^2	0.888
Dep. var. mean	18.2

Clustered (Region Pair) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

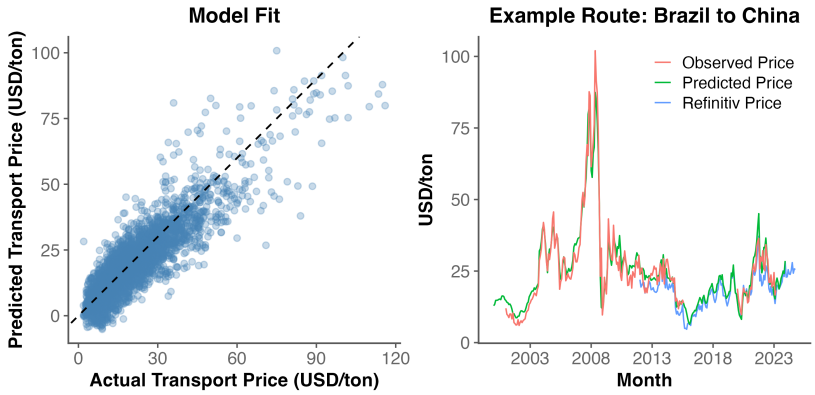


Figure 21: Predictive Regression for Shipping Prices p_T. The left panel reports the estimated coefficients from the regression used to predict shipping prices. The dependent variable is the observed trip-level fixture price; the independent variables include distance traveled, fuel costs, the Baltic Dry Index (BDI), and region fixed effects. The right panel shows the fit of the regression in aggregate as well as one route as an example.The key takeaway is the regression's high predictive power.

Commodity Price Inflation, $\log\left(P_{ijct}/P_{ijc,t-4}\right)$**.** We define importer prices as $P_{ijct} = \text{ProducerPrice}_{ijct} + \text{TransportPrice}_{ijct}$, corresponding to $p_x + p_T$.

Total Time at Port, TT_{ijct}**.** Our analysis focuses on time at port at exporting ports. We first restrict our sample of port calls to ports that have ever traded at least one of the major commodities in our sample. Recall that our port call data measures both the port service time and wait time at anchorage for every port call. However, before we aggregate, it is important to account for heterogeneity in ship size, which affects the ship's service time. To this end, we construct a size-adjusted measure of service time, T, by residualizing observed service time with respect to vessel size. Specifically, we estimate

$$T_{sp} = \beta \log\left(\text{Vessel Size}_s\right) + \gamma_p + \epsilon_{sp}$$

where s index port calls, p index ports. We then define residualized service time as the fitted service time for a Panamax vessel (which is approximately the median size vessel) at port p, plus the residual for port call s (i.e. $\hat{T}_{sp} = \hat{\beta} \log\left(\text{Panamax}\right) + \hat{\gamma_p} + \hat{\epsilon_{sp}}$).

The sum of this residualized service time and the wait time is our measure of total time spent at

a port, TT_{pt}. To aggregate TT_{pt} to the importer–exporter–commodity–time, or $i-j-c-t$, level, we compute the share of each exporting port p in shipments of commodity c from exporter j to importer i.[28] We then define $TT_{ijct} = \sum_{p \in j} w_{p,ijc} TT_{pt}$, where $w_{p,ijc}$ denotes port p's market share in exporter j's shipments of commodity c to importer i.

B.2 Price Regressions

In Section 2.3 we IV regress commodity prices on total time at port, TT_{ijct}. Our instrument for TT_{ijct}, which we denote here by $Shock_{ijct}$, is meant to capture variation in the port service time, and hence total time at port, that is driven by unexpected disruptions. To construct it, we regress service time on a set of observable determinants of port operations and define the residual from this regression as our measure of an unexpected service-time shock. The set of observables includes log ship size, the maximum number of ships that are serviced at the same time, port interacted with quarter-of-year fixed effects, month fixed effects, and commodity fixed effects. This instrument captures disruptions in port service operations within a month, that we then aggregate at the quarterly level for every importer-exporter-commodity pair, based off each port's relative trade share, following what we did for TT_{ijct}. These disruptions increase total time at port and, through the increased time at port they affect demand, since shippers prefer ports with shorter time at port. However, these disruptions do not affect demand directly.

Table 3 reports the results from the first stage discussed above. Table 4 provides a validation of the instrument, by regressing it on known disruptions, such as localized weather (controlling for weather in the broader area), as well as whether the arrival of the ship happened during the weekend.

[28] We drop observations that exceed the 99th percentile of residualized service time.

Model:	Total time at port (1)
Unex. disruption at port	0.8157***
	(0.0666)
importer-exporter	Yes
commodity-quarter	Yes
importer-quarter	Yes
exporter-quarter	Yes
R^2	0.80
Observations	60,983
KP F Stat	150.0

Table 3: Impact of port congestion on price changes: First-Stage. This is the first-stage regression for equation 1. The dependent variable is total time at port. The independent variables include combinations of our instrument, importer-exporter fixed effects, commodity-quarter fixed effects, importer-quarter fixed effects and exporter-quarter fixed effects. The instrument is constructed from unexpected disruptions in port service time.

Model:	Unexpected Disruption at Port (1)
1(service starts on a weekend)	0.083
	(0.136)
Country average rain level	1.31***
	(0.175)
Local mean rain level (10km)	0.833***
	(0.042)
Port-Month	Yes
Port-Commodity	Yes
R^2	0.16
Observations	2,069,019

Table 4: Instrument Validation. The unit of observation is a port call. The dependent variable is our measure of unexpected disruption (in hours) at the port-call level. The independent variables include an indicator for whether the port call begins service on a weekend and the average local precipitation level - measured in millimeters per day - within a 10 km radius of the port centroid during the port call. We control for the average precipitation level in the country during the same port-call period. We also control for port-month fixed effects and port-commodity fixed effects. Precipitation data are from the NASA Goddard Earth Sciences Data and Information Services Center (GES DISC); we use the product "GPM IMERG Final Precipitation L3 1 day 0.1 degree x 0.1 degree V07" (Huffman et al., 2023), accessed here.

	OLS	IV
Model:	(1)	(2)
Constant	-0.0754***	
	(0.0230)	
TT at port	0.0132***	-0.0449**
	(0.0036)	(0.0207)
importer-exporter		Yes
commodity-quarter		Yes
importer-quarter		Yes
exporter-quarter		Yes
R^2	0.0002	0.15
Observations	65,631	64,712
KP F Stat, TT at port		168.5

Table 5: Impact of port congestion on quantity traded. Regression results for equation (1) where the dependent variable is the year-over-year log change in quantities traded instead of the price. The independent variables include total time at port, importer-exporter fixed effects, commodity-quarter fixed effects, importer-quarter fixed effects and exporter-quarter fixed effects. The first column reports OLS estimates. The second column reports the IV estimate, where we use unexpected disruption as the instrument. Standard errors are clustered at importer-exporter level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

C Proofs

C.1 Proof of Proposition 1

Let

$$\bar{Q}_{\text{ship}} = \sup\left\{q \in \mathcal{R}^+ : f(q) > 0\right\}$$

where

$$f(q) \equiv \frac{S}{2d + T + W_{\text{port}}\left(q, S, d, K, T\right)} - q$$

We want to show that $\bar{Q}_{\text{ship}} = \frac{S}{2d+T+W_{\text{port}}\left(\bar{Q}_{\text{ship}},S,d,K,T\right)}$, or $f\left(\bar{Q}_{\text{ship}}\right) = 0$.

First, it is easy to see that this set is non-empty; indeed it includes at least zero, since $f(0) = \frac{S}{2d+T} > 0$. Moreover, note that the wait time for the ships is bounded above since the number of ships in the system is finite. Therefore, the supremum exists in $\mathbb{R}$.

Since $W_{\text{port}}(\cdot)$ is continuous in q, so that f is also continuous in q. First, we show that $f(\bar{Q}_{\text{ship}}) \geq 0$. By definition of sup, there exists a sequence (q_n) in the set $\{q \in D : f(q) > 0\}$ converging to $\bar{Q}_{\text{ship}}$. By continuity of f, $f(q_n) \to f\left(\bar{Q}_{\text{ship}}\right)$. Since $f(q_n) > 0$, $f\left(\bar{Q}_{\text{ship}}\right) \geq 0$.

Now we want to show equality. Assume by contradiction that $f(\bar{Q}_{\text{ship}}) > 0$. Then by continuity of f and that $\bar{Q}_{\text{ship}}$ is bounded above (denote the upper bar by $\bar{q}$) , there exists a small enough ϵ such that $f(q) > 0$ for all $q \in \left(\bar{Q}_{\text{ship}} - \epsilon, \bar{Q}_{\text{ship}} + \epsilon\right)$, and $\bar{Q}_{\text{ship}} + \epsilon < \bar{q}$, where , $\bar{q}$, is the upper bound. Now take $\hat{q} = \bar{Q}_{\text{ship}} + \frac{\epsilon}{2}$. Note that this contradicts the fact that $\bar{Q}_{\text{ship}}$ is the supremum. Hence the strict inequality must be false.

C.2 Proof of Lemma 3

We show Lemma 3 holds for $z \in \{S, K\}$; the other two can be shown analogously. When the ship constraint binds, in steady state we have that

$$Q = \bar{Q}^{max} = \frac{S}{2d + T + W_{\text{port}}\left(Q, S, d, K, T\right)}$$

We begin by calculating dQ/dS. Let us define

$$F(Q, S) = \frac{S}{2d + T + W_{\text{port}}\left(Q, S, d, K, T\right)} - Q$$

Then, by the implicit function theorem, we have,

$$\begin{aligned}\frac{dQ}{dS} &= -\frac{F_S}{F_Q} \\ &= \frac{2d + T + W_{\text{port}}(Q, S, d, K, T) - S\frac{\partial W_{\text{port}}(Q,S,d,K,T)}{\partial S}}{(2d + T + W_{\text{port}}(Q, S, d, K, T))^2 + S\frac{\partial W_{\text{port}}(Q,S,d,K,T)}{\partial Q}}\end{aligned}$$

Note however that $\frac{\partial W_{\text{port}}(Q,S,d,K,T)}{\partial S} = 0$; this is because the port wait time is a function of the arrival rate at the port, denote it by λ, as well as the port service rate T, and port infrastructure K. The arrival rate λ is the departure rate from the Ship Queue and thus $\lambda(q, S, d, K, T)$. Hence the fleet S affects W_{port} only through λ and not directly. Finally, since $\frac{\partial W_{\text{port}}}{\partial Q} > 0$, we get $\frac{dQ}{dS} > 0$ and the result is obtained.

We now turn to dQ/dK. Similarly, define

$$F(Q, K) = \frac{S}{2d + T + W_{\text{port}}(Q, S, d, K, T)} - Q$$

so that

$$\begin{aligned}\frac{dQ}{dK} &= -\frac{F_K}{F_Q} \\ &= -\frac{S}{S + (2d + T + W_{\text{port}}(Q, S, d, K, T))\frac{\partial W_{\text{port}}(Q,S,d,K,T)}{\partial K}\frac{\frac{\partial W_{\text{port}}(Q,S,d,K,T)}{\partial Q}}{\frac{\partial W_{\text{port}}(Q,S,d,K,T)}{\partial K}}}\end{aligned}$$

Assuming the intuitive property that $\frac{\partial W_{\text{port}}(Q,S,d,K,T)}{\partial K} < 0$ proves the result.

D Empirical Implementation, Details

In this Appendix we discuss details related to the calibration of the model presented in Section 4. As above, we focus on a subset of bulk commodities: cement, ore, coal, fertilizer, soybeans, wheat, and corn. We also exclude domestic trade flows, defined as shipments for which the loading port and discharging port are located in the same country.

Quantity, Q. For each quarter t, we calculate total world trade volume by summing the tonnage recorded in the AIS data across all shipments. We convert this volume into an implied number of shipments by dividing total tonnage by the size (DWT) of the median vessel in our data. Thus, one shipment corresponds to the cargo capacity of a median ship.

In the calibration exercise for 2004-2009, since AIS data are not available (they do not reliably exist prior to 2013), we rely on quantities reported in COMTRADE. A key challenge is that quarterly COMTRADE coverage is incomplete in the early 2000s. To address this, we use annual COMTRADE totals for aggregate quantities and allocate them to quarters based on the average quarterly distribution of trade across quarters from the subset of countries that report this data. Finally, we reconcile the level difference between COMTRADE and AXS by scaling COMTRADE quantities using their relative ratio in the first half of 2013.

Port Infrastructure, K. We first compute a measure of infrastructure $K_{i,\tau}$ for every port i, in every period τ, where $\tau \in \{\text{2013-2016, 2017-2019, 2020-2023}\}$. For a given port i, we calculate, for each day, the maximum total tonnage of ships being serviced simultaneously, conditional on at least one additional vessel waiting for service. For each period τ, we then define port infrastructure as the 75th percentile of this daily measure. Port capacity is therefore measured in deadweight tons. This approach is similar to Brancaccio et al. [2024]. We next convert port capacity into median-ship equivalents by normalizing by the size of the median vessel.

Note that in our calibrated model, however, capacity is not a port-level object: the model features a single exporting country, so the relevant notion of capacity is the effective aggregate capacity of all ports in exporting countries, rather than the physical sum of all port capacities. We therefore aggregate the port-level measures $K_{i,\tau}$ across all ports in exporting countries.

To this end, we first compute world $\tilde{K}$ for the last period in our sample, 2020-2023. Let $|\mathcal{P}|$ denote the set of world ports. We define $\tilde{K}$ of the aggregate exporting port such that

$$\tilde{K} \times \text{Weighted shipment per berth} = \text{Total Export Shipment}, \tag{11}$$

where

$$\text{Weighted shipment per berth} = \sum_{j \in |\mathcal{P}|} \frac{\text{ExportingShipments}_j}{\sum_{j' \in |\mathcal{P}|} \text{ExportingShipments}_{j'}} \times \text{Shipment per berth}_j. \tag{12}$$

In this way, the measure defined in (12) computes the average number of shipments handled by each berth, weighting each port by its share of the world's exports. Correspondingly, $\tilde{K}$ is the level of capacity of a single aggregate exporting port that would reproduce total observed export shipments if it had this average shipments-per-berth ratio. At a high level, instead of assuming exporting shipments are

distributed evenly across all berths in the world, we are adjusting the global $\tilde{K}$ to reflect where exporting activities are actually directed. [29]

To recover capacity in 2013-2016 and 2017-2019, we project $\tilde{K}$ back, proportionally by the total stock of world port K.

For the calibration for the 2004-2009 period, AIS data are not available (since we do not have them prior to 2013). We therefore calibrate port K to its average level in 2013. This likely provides an upper bound for the true value K and therefore we are less likely to overestimate the importance of port bottlenecks.

Service time, T. We calculate monthly service time at each port as described in Section B above. We then aggregate these monthly service times across ports worldwide to obtain service time for our representative exporting country, using each port's average share of world exports as weight. Finally, we aggregate the monthly series to the quarterly level using the total number of trips in each month as the relevant weight.

For the calibration for the 2004-2009 period, AIS data are not available (since we do not have them prior to 2013). We therefore calibrate service time T to its average level in 2013. This likely provides a conservative estimate of T.

Marginal Cost of a Trip, c. The monthly marginal cost of a trip includes three components: fuel costs, port charges and a fixed operating cost. The first component, fuel cost, is equal to the average distance travelled (two voyage legs of 2.9 weeks each) times monthly fuel costs.

Port charges are set to USD 3 per ton (Brancaccio et al., 2024). Because port charges vary little over time, we treat them as time invariant.

Finally, we calibrate the fixed operating cost of a voyage using timecharter rates as a proxy for the daily fixed cost of operating the vessel when prices are low. Under a time-charter contract, the shipowner is responsible for the vessel's operating costs, while the charterer covers voyage-specific expenses such as fuel and port charges. To recover the vessel's operating costs, we assume that at its lowest level, observed time-charter rates allow the shipowner to only cover these costs, consistent with our model. Therefore, we calibrate the daily vessel's operating costs to coincide with the the first percentile of the observed time-charter rate distribution in our sample period. Because time-charter rates are quoted in U.S. dollars

[29] Note that if the export shipment-per-berth ratio is homogenous across all ports, then our measure of $\tilde{K}$collapses to the simple sum of ports' capacity. Conversely, if a single port handles all export activities in the world, then $\tilde{K}$ collapses to the capacity of that particular port.

per day, we convert them into voyage-level fixed costs by multiplying the daily rate by the calibrated trip duration, which we set equal to 2×2.9 weeks.

Prices, p, p_T, p_x. As discussed in Appendix B above, prices are measured at the exporter– importer– commodity– quarter level. To aggregate them to the world level, we proceed in two steps. First, for each commodity and quarter, we compute a weighted average across exporter–importer pairs, using each pair's traded quantity as the weight. To avoid introducing time-varying composition effects, we measure these weights using the median traded quantity for each $i - j$ pair over time. Second, we aggregate from the commodity-quarter level to the quarterly level by taking a weighted average across commodities, where the weights are each commodity's median traded quantity over time.[30]

Fleet Size, S. We observe the world total fleet, measured in DWT, at a monthly frequency from Clarkson data. We first collapse this data to semester level. To account for domestic trade and trade in commodities outside our sample, we scale this measure by the share of cross-border trade attributable to the commodities included in our sample.

Transportation Capacity, $Q^{\max}$. To compute $Q^{\max}$ we need the minimum of K/T and $S/\left(2d + T + W_{\text{port}}\right)$. Although the first term is directly computable from the data, the second requires us to simulate the model in order to obtain the last term, W_{port}. For this reason, we directly simulate $Q^{\max}$ as follows. We simulate the sequence of queues for every quarter by setting the arrival rate of exporters high enough, so that ships effectively never wait for shipments to arrive. In other words, we want to ensure that the system is constrained and the equilibrium quantity transported is given by $Q^{\max}$. We compute the realized throughput per unit of time after discarding a warm-up window, and define $Q^{\max}$ as this realized throughput. Then, by comparing $Q^{\max}$ and K/T we can distinguish the role of the port constraint and ship constraint in determining $Q^{\max}$.

Next, we create bootstrap samples for $Q^{\max}$ between 2013 and 2023, in Figure 9, as follows. First, note that throughout 2015 to 2023, the port constraint pins down $Q^{\max}$, because $Q^{\max} = K/T$. With that, for each port, we draw port calls with replacement within each quarter to obtain bootstrap estimates of service times, i.e. T, for each port-quarter. We similarly resample days with replacement to obtain bootstrap estimates of the maximum number of ships serviced simultaneously, i.e. K, for each port. Port-level measures of T and K are then aggregated to the world level using the procedure described

[30] To maintain consistency with our measure of Q, all prices are expressed in dollars per shipment before being fed into the model.

earlier in this Appendix. We repeat this procedure 100 times. Second, note that the ship constraint plays a more prominent role from 2013 to 2014, therefore variation in T and K affects $Q^{\max}$ through the ship constraint as well. We therefore re-run the simulation described in the previous paragraph for each bootstrap sample of T and K, and use the resulting simulated capacity as the corresponding bootstrap draw of $Q^{\max}$.

References

George Alessandria, Shafaat Yar Khan, Armen Khederlarian, Carter Mix, and Kim J. Ruhl. The aggregage effects of global and local supply chain disruptions: 2020-2022. *Journal of International Economics*, 146(103788), 2023.

Treb Allen and Costas Arkolakis. Trade and the topography of the spatial economy. *Quarterly Journal of Economics*, 129(3):1085–1140, 2014.

Treb Allen and Costas Arkolakis. The welfare effects of transportation infrastructure improvements. *The Review of Economic Studies*, 89(6):2911–2957, 2022.

Mary Amiti, Oleg Itskhoki, and David E. Weinstein. What drives us import price inflation? *AEA Papers and Proceedings*, 114:106–111, 2024.

James E. Anderson and Eric Van Wincoop. Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1):170–192, 2003.

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

AXSMarine. Nearly 300 Less Ships Per Week Through Suez Canal Due To Red Sea Crisis, February 2024. URL `https://public.axsmarine.com/blog/nearly-300-less-ships-through-suez-canal-through-red-sea-crisis`.

Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti. The causal effects of global supply chain disruptions on macroeconomic outcomes: evidence and theory. *NBER Working Paper*, 32098, 2024.

Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li, Ricardo Marto, and Francesco Zanetti. Shipping to America. *Manuscript*, June 2026a.

Xiwen Bai, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti. State Dependence of Monetary Policy During Global Supply Chain Disruptions, May 2026b. URL `https://www.nber.org/papers/w35209`.

Samuel Bailey. Competition and coordination in infrastructure: Port authorities' response to the panama canal expansion. *Unpublished Manuscript*, 2021.

Gianluca Benigno, Julian di Giovanni, Jan J.J. Groen, and Adam I. Noble. The gscpi: A new barometer of global supply chain pressures. *Federal Reserve Bank of New York Staff Reports*, (1017), 2022.

Ben Bernanke and Olivier Blanchard. What Caused the US Pandemic-Era Inflation? *American Economic Journal: Macroeconomics*, 17(3):1–35, July 2025. ISSN 1945-7707. doi: 10.1257/mac.20230195. URL https://www.aeaweb.org/articles?id=10.1257/mac.20230195.

Joaquin Blaum, Federico Esposito, and Sebastian Heise. Input Sourcing under Climate Risk: Evidence from U.S. Manufacturing Firms. *The Review of Economic Studies*, May 2026.

Giulia Brancaccio, Myrto Kalouptsidi, and Theodore Papageorgiou. Geography, transportation, and endogenous trade costs. *Econometrica*, 88(2):657–691, 2020.

Giulia Brancaccio, Myrto Kalouptsidi, Theodore Papageorgiou, and Nicola Rosaia. Search frictions and efficiency in decentralized transport markets. *Quarterly Journal of Economics*, 138(4):2451–2503, 2023.

Giulia Brancaccio, Myrto Kalouptsidi, and Theodore Papageorgiou. Investment in infrastructure and trade: The case of ports. Technical report, 2024.

Giulia Brancaccio, Myrto Kalouptsidi, and Theodore Papageorgiou. Rigidities in transportation and supply chain disruptions. *AEA Papers and Proceedings*, 115:543â50, May 2025. doi: 10.1257/pandp.20251079. URL https://www.aeaweb.org/articles?id=10.1257/pandp.20251079.

Christian Broda, Nuno Limao, and David E Weinstein. Optimal tariffs and market power: the evidence. *American Economic Review*, 98(5):2032–2065, 2008.

Leah Brooks, Nicolas Gendron-Carrier, and Gisela Rua. The local impact of containerization. *Journal of Urban Economics*, 126(103388), 2021.

Yann Calvo Lopez and Benjamin Golub. Comment on Hidden Exposure: Measuring US Supply Chain Reliance. *Brookings Papers on Economic Activity*, 2023(2):146–167, September 2023. ISSN 1533-4465. doi: 10.1353/eca.2023.a935416. URL https://muse.jhu.edu/article/935416.

Yan Carrière-Swallow, Pragyan Deb, Davide Furceri, Daniel Jiménez, and Jonathan D. Ostry. Shipping costs and inflation. *IMF Working Papers*, 2022.

Vasco M Carvalho, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi. Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321, 2021.

Juanma Castro-Vincenzi. Climate Hazards and Resilience in the Global Car Industry. *Working Paper*, June 2024.

Diego Comin, Robert C. Johnson, and Callum Jones. Supply chain constraints and inflation. *NBER Working Paper*, 2023.

Diego Comin, Robert C. Johnson, and Callum Jones. Import constraints. *AEA Papers and Proceedings*, 114: 118–123, 2024.

Dave Donaldson. Railroads of the raj: Estimating the impact of transportation infrastructure. *forthcoming, American Economic Review*, 2012.

Dave Donaldson and Richard Hornbeck. Railroads and american economic growth: A "market access" approach. *Quarterly Journal of Economics*, 131(2):799–858, 2016.

Cesar Ducruet, Reka Juhasz, David Krisztian Nagy, and Claudia Steinwender. All aboard: The effects of port development. *Journal of International Economics*, 151:103963, 2024.

Jason Dunn and Fernando Leibovici. Navigating the waves of global shipping: Drivers and aggregate implications. *mimeo, Federal Reserve Bank of St. Louis*, 2025.

Jonathan Eaton and Samuel Kortum. Technology, geography and trade. *Econometrica*, 70(5):1741–1779, 2002.

Matthew Elliott, Benjamin Golub, and Matthew V. Leduc. Supply Network Formation and Fragility. *American Economic Review*, 112(8):2701–2747, August 2022.

Pablo D. Fajgelbaum and Edouard Schaal. Optimal transport networks in spatial equilibrium. *Econometrica*, 88 (4):1411–1452, 2020.

Simon Fuchs and Woan Foong Wong. Multimodal transport networks. *NBER Working Paper*, 35065, 2026.

Simon Fuchs, Fernando Leibovici, and Woan Foong Wong. Geopolitical risk and foreign dependence in global shipping. 2026.

Sharat Ganapati, Woan Foong Wong, and Oren Ziv. Entrepot: Hubs, scale, and trade costs. *American Economic Journal: Macroeconomics*, 16(4):239–78, October 2024.

Ezequiel Garcia-Lembergman, Natalia Ramondo, Andres Rodriguez-Clare, and Joe Shapiro. Carbon emissions in the global economy. *mimeo, University of Berkeley*, 2025.

Gita Gopinath, Pierre-Olivier Gourinchas, Andrea F. Presbitero, and Petia Topalova. Changing global linkages: A new Cold War? *Journal of International Economics*, 153:104042, January 2025.

Gene M Grossman, Elhanan Helpman, and Hugo Lhuillier. Supply chain resilience: Should policy promote diversification or reshoring? *NBER Working Paper*, 29330, 2021.

Inga Heiland, Andreas Moxnes, Karen Helene Ulltveit-Moe, and Yuan Zi. Trade from space: Shipping networks and the global implications of local shocks. *The Review of Economics and Statistics*, pages 1–45, 07 2025.

G. J. Huffman, E. F. Stocker, D. T. Bolvin, E. J. Nelkin, and J. Tan. Gpm imerg final precipitation l3 1 day 0.1 degree x 0.1 degree v07, 2023. Accessed: [Data Access Date].

David Hummels. Transportation costs and international trade in the second era of globalization. *The Journal of Economic Perspectives*, 21(3):131–154, 2007.

David Hummels and Sasha Skiba. Shipping the good apples out? an empirical confirmation of the alchian-allen conjecture. *Journal of Political Economy*, 112(6):1384–1402, 2004.

David Hummels, Volodymyr Lugovskyy, and Sasha Skiba. The trade reducing effects of market power in international shipping. *Journal of Development Economics*, 89(1):84–97, 2009.

Jota Ishikawa and Nori Tarui. Backfiring with backhaul problems: Trade and industrial policies with endogenous transport costs. *Journal of International Economics*, 111:81–98, 2018.

Cleve Jones and Robert Wright. The mounting strains on global shipping. *Financial Times*, May 2024. URL `https://www.ft.com/content/a03da8f6-b468-4a86-8db5-83838f7d6409?syn-25a6b1a6=1`.

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

Diego R. Kanzig and Ramya Raghavan. Supply chain shocks and the macroeconomy: Evidence from global shipping disruptions. *mimeo, Northwestern University*, 2026.

Benny Kleinman, Ernest Liu, Stephen J. Redding, and David Xu. International Trade in an Uncertain World, July 2026. URL `https://www.nber.org/papers/w35416`.

Thomas Kohler, Jean-Paul L'Huillier, Gregory Phelan, and Maximilian Weiss. Why Do Supply Disruptions Lead to Inflation (While Demand Booms Do Not)? Survey Evidence from the COVID Pandemic. *CESifo Working Papers*, February 2026.

Vasily Korovkin, Alexey Makarin, and Yuhei Miyauchi. Supply Chain Disruption and Reorganization: Theory and Evidence From Ukraineâs War. *The Review of Economic Studies*, 93(4):2750–2783, July 2026. ISSN 0034-6527. doi: 10.1093/restud/rdaf080. URL `https://doi.org/10.1093/restud/rdaf080`.

Richard C. Larson and Amedeo R. Odoni. *Urban Operations Research*. Dynamic Ideas, 2nd edition, 2007.

Marc J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.

Nektarios A. Michail, Konstantinos D. Melas, and Lena Cleanthous. The relationship between shipping freight rates and inflation in the euro area. *International Economics*, 172:40–49, 2022.

Stephen J. Redding. Goods trade, factor mobility and welfare. *Journal of International Economics*, 101:148–167, 2016. ISSN 0022-1996. doi: https://doi.org/10.1016/j.jinteco.2016.04.003. URL https://www.sciencedirect.com/science/article/pii/S0022199616300502.

Stephen J Redding. Evaluating transport improvements in spatial equilibrium. *NBER Working Paper*, 33532, 2025.

Stephen J. Redding and Matthew A. Turner. Chapter 20 - transportation costs and the spatial organization of economic activity. In Gilles Duranton, J. Vernon Henderson, and William C. Strange, editors, *Handbook of Regional and Urban Economics*, volume 5 of *Handbook of Regional and Urban Economics*, pages 1339–1398. Elsevier, 2015.

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

The Economist. A new Suez crisis threatens the world economy. *The Economist*, December 2023. ISSN 0013-0613. URL https://www.economist.com/international/2023/12/16/a-new-suez-crisis-threatens-the-world-economy.

Ward Whitt. The queueing network analyzer. *Bell System Technical Journal*, 62(9):2779–2815, November 1983. doi: 10.1002/j.1538-7305.1983.tb03204.x.

Woan Foong Wong. The round trip effect: Endogenous transport costs and international trade. *American Economic Journal: Applied Economics*, 14(4):127–166, 2022.