NBER WORKING PAPER SERIES

ALL ABOARD: THE EFFECTS OF PORT DEVELOPMENT

César Ducruet
Réka Juhász
Dávid Krisztián Nagy
Claudia Steinwender

Working Paper 28148
http://www.nber.org/papers/w28148

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge. MA 02138
November 2020, Revised April 2023

We thank Treb Allen, David Atkin, Leah Brooks, Arnaud Costinot, Don Davis, Dave Donaldson, Joseph Doyle, Costas Arkolakis, Nicolas Gendron-Carrier, Matt Grant, Gordon Hanson, Tarek Hassan, Tom Holmes, David Hummels, Amit Khandelwal, Giampaolo Lecce, Nelson Lind, Nina Pavcnik, Giacomo Ponzetto, Jim Rauch, Steve Redding, Roberto Rigobon, Andrés Rodríguez-Clare, Esteban Rossi-Hansberg, Daniel Sturm, Tavneet Suri, Jaume Ventura, Jon Vogel and David Weinstein for helpful comments and discussions. We thank Bruce Blonigen, Mario Martín Antón and their co-authors for kindly sharing data. We thank staff at various port authorities for helping with our information requests. Olalekan Bello, Sabrina Chen, Naman Garg, Yi Jie Gwee, Hamza Husain, Felix Iglhaut, Rodrigo Martínez Mazza, Emanuela Migliaccio, Véronica Perez, Shuhua Si, Yue Yu, Howard Zihao Zhang and a team of Columbia University undergraduate students provided outstanding research assistance. César gratefully acknowledges financial support from the ERC (Starting Grant 'World Seastems' No. 313847). Réka is grateful for financial support from the Provost's Office at Columbia. Dávid acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019- 000915-S and SEV-2015-0563) and through a Juan de la Cierva Grant (FJCI-2017-34728), the Spanish State Research Agency and the European Social Fund (RYC2019-027620-I / AEI / 10.13039/501100011033) and from the Generalitat de Catalunya, through CERCA and SGR Programme (2017-SGR-1393). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

All Aboard: The Effects of Port Development
César Ducruet, Réka Juhász, Dávid Krisztián Nagy, and Claudia Steinwender
NBER Working Paper No. 28148
November 2020, Revised April 2023
JEL No. F6,O33,R40

ABSTRACT

Seaports facilitate the fast flow of goods across space, but ports also entail local costs borne by host cities. We use the introduction of containerized shipping to explore the effects of port development. At the local level, we find that seaport development increases city population by making a city more attractive, but this market access effect is offset by costs which make the city less attractive. At the aggregate level, we find that the local costs associated with port development are heterogeneous across cities and reduce aggregate welfare gains, which however are still positive and substantial.

César Ducruet
French National Centre for
Scientific Research (CNRS)
13 rue du Four
PARIS 75006
France
cdu@parisgeo.cnrs.fr

Réka Juhász
Vancouver School of Economics
University of British Colunbia
6000 Iona Driv
Vancouver, BC V6T 1L4
and NBER
reka.juhasz@ubc.ca

Dávid Krisztián Nagy
CREI
Ramon Trias Fargas, 25-27
08005 Barcelona
Spain
dnagy@crei.cat

Claudia Steinwender
Ludwig Maximilian University of Munich
Department of Economics
Ludwigstr. 28 (Vgb.)
Munich D-80539
Germany
claudia.steinwender@econ.lmu.de

An online appendix is available at http://www.nber.org/data-appendix/w28148

> Across the planet, the expansion of seaports is becoming tougher (...). Space in the right locations is scarce.
>
> *The Economist*, January 14th, 2023

Introduction

Seaports play a vital role in the global trading system, handling over 80% of world merchandise trade in 2018 in terms of volume (UNCTAD, 2019). Efficient, modern facilities that provide ample space for the fast loading and unloading of containers are a precondition for a country to participate in global production networks (Rodrigue, 2016, p. 131). Despite their importance, little is known about the economic effects of ports. What determines which coastal cities become important ports? What are the aggregate gains from port development? Which cities reap the benefits, and which pay the costs of port development?

In this paper, we study these questions by exploiting a major technological shock to port development: containerization, that is, the handling of cargo in standardized boxes. Our analysis sheds light on a novel mechanism that affects i) the economic geography of ports, ii) the gains from port development, and iii) the distribution of these gains. This mechanism is driven by the *local costs of port development*.

Modern port development entails at least two costs that are borne by host cities. First, ports occupy large amounts of land in their host cities. For example, the ports of Antwerpen and Rotterdam occupy more than 30% of the metropolitan area of the city, while in Los Angeles 85% of total truck traffic is accounted for by port truck traffic on some highway segments (OECD, 2014, p. 17). The costs associated with space have become particularly salient with recent supply chain disruptions. The overflow of containers in major ports such as Long Beach, California that did not have slack capacity highlights the extent to which many modern ports are space-constrained.[1]

Second, ports may induce large-scale local disamenities such as noise and pollution. In Hong Kong, more than half of the sulphur dioxide emissions are related to shipping (OECD, 2014, p. 17). As a recent article in The Economist (2023) highlights, the localized land and environmental costs of port development are arguably some of the most pressing current challenges for ports.

In the first part of our analysis, we assemble a unique panel dataset of city populations and shipping flows to document the local effects of port development across the globe. We use the introduction of containerized shipping to explore these effects. To isolate exogenous variation in a city's suitability for containerization, we build on a previous literature that has shown that access to deep water at the port is important for containerization (Brooks, Gendron-Carrier, and Rua, 2021; Altomonte, Colantone, and Bonacorsi, 2018). We construct a novel measure of 'naturally

[1] E.g., https://qz.com/2079345/cargo-ships-containers-are-piling-up-in-long-beach.

endowed' depth (as distinct from depth attained by dredging) based on granular data on oceanic depths around a port.

Using this exogenous measure of suitability to containerization, we document three empirical facts. First, we show that cities exogenously more suited to containerization witnessed a boom in shipping flows after the onset of containerization. This fact suggests that containerization increased these cities' market access by lowering their shipping costs. Second, we find that the shipping boom was less pronounced in cities where land is scarce due to geographic constraints. This fact reflects the importance of land costs for port development. Third, we show that the increase in local shipping did not translate into population inflows for the average city: our IV estimates show an effect of increased shipping on population growth that is both economically and statistically insignificant. This fact suggests that the local costs of port development from land or other sources can fully offset the benefits from better market access. The economic geography literature has traditionally focused on only these market access benefits (Donaldson and Hornbeck, 2016; Redding and Turner, 2015).

In the second part of the paper, we develop a general equilibrium model that can be used to quantify the aggregate and distributional impacts of port development. The model adds an endogenous port development decision to an otherwise standard economic geography model of trading cities. Port development is costly for two reasons: it requires scarce local land, and it creates disamenities in the port city. As a result, the model incorporates not only the standard market access effect, but also both types of local port development cost suggested by the qualitative literature and our empirical facts. Whether a city ultimately gains in population is the outcome of the trade-off between the market access benefits and the local (land use and disamenity) costs of port development.

The model generates an empirically testable relationship to disentangle the local benefits and costs of port development. More specifically, it predicts that the benefits of port development on city population are captured by market access, while the costs of port development are captured by shipping, once market access is controlled for. Using port depth and changing city amenities to generate exogenous variation in shipping and market access, we confirm a *negative* causal effect of shipping on city population once the *positive* causal effect of market access is controlled for.

Finally, we quantify the aggregate and distributional effects of port development by taking the model to the data. We use data on shipping flows, city GDP and population in 1990 to back out cities' unobserved model fundamentals. Armed with these fundamentals, we conduct two counterfactual simulations to shed light on the importance of the local costs of port development.

In our first counterfactual, we simulate the pre-containerization equilibrium in the model by *undoing* the containerization shock. We test the model by showing that it can successfully replicate the three empirical facts: deeper port cities see faster shipping growth; this is less pronounced in cities where land is scarce; and shipping growth does not translate into growth in city population. We also show that a traditional ("benchmark") model of trade and geography, in which transport cost reductions are *exogenous* and *costless*, fails to replicate these facts. First, the benchmark model does not predict a significant relationship between land scarcity and shipping growth. Second, it predicts that port city populations grow significantly as a result of containerization, unlike in the data. Additionally, while we find a quantitatively meaningful role for the disamenity effects of port development, we show that they alone (without land use costs) cannot explain our three empirical facts.

Our estimates suggest that containerization increased world welfare by 3.23%, the ratio of world trade to GDP by 4.1 percentage points, and the median port size relative to city area by 5 percentage points. In a model-based decomposition, we find that the aggregate resource cost of increased land use amounted to 0.64% of world GDP, reducing the welfare gains from containerization by 18%. This result highlights that the local costs of port development are important not only for where port activity is located, but also for how much the world as a whole gained from containerization. We also find an additional welfare gain from cities' endogenous specialization in port- and non-port activities, depending on their comparative advantage. These specialization gains offset 63% of the resource costs of containerization, but they do not compensate for all the costs.

In our second counterfactual, we examine the effects of targeted port development policies. We focus on a setting similar to the 'Maritime Silk Road' project – a large set of port investments undertaken by China in South Asian, African and European ports. Our findings suggest that targeted port development has the potential for large distributional effects triggered by the reallocation of shipping activity. The model predicts a large decline in shipping in Singapore (a non-targeted port which we estimate to lose about 50% of its shipping flows), which is driven by the fact that shipping activity reallocates to nearby, targeted ports. Crucially, this initial shock is amplified by less endogenous port development in Singapore as demand for port services falls. However, despite losing a sizeable fraction of its shipping flows, Singapore gains 1% in GDP, as resources reallocate to Singapore's highly productive non-port activities. This illustrates that, because of the costs of port development, gains and losses in shipping do not translate directly into gains in real GDP. These findings highlight the importance of accounting for the endogenous port development mechanism when quantifying how the gains from targeted port development are distributed across space. More speculatively, they question the wisdom of highly productive, expensive cities such as Hong Kong and Singapore continuing to specialize heavily in port services.

Related literature. A recent, growing literature provides evidence that better trading opportunities lead to local benefits through increasing market access (Donaldson and Hornbeck, 2016; Redding and Turner, 2015), which may induce city development (Bleakley and Lin, 2012; Armenter, Koren, and Nagy, 2014; Nagy, 2022). Some of these studies focus on city development at port locations in particular (Fujita and Mori, 1996; Coşar and Fajgelbaum, 2016; Fajgelbaum and Redding, 2022). We contribute to this literature by showing that trade-induced development can also have substantial local costs. While the potential for transport infrastructure to put a strain on scarce local resources has long been recognized theoretically (Solow and Vickrey, 1971; Solow, 1972; Pines and Sadka, 1985), the effect has not been estimated empirically. This mechanism also relates the paper to the 'Dutch disease' literature, which shows that booming industries can entail significant costs through competing with other (tradable) sectors for local resources (Corden and Neary, 1982; Krugman, 1987; Allcott and Keniston, 2017).[2] Relative to this literature, our setting contains the potential for not only costs but also gains, as booming port activities benefit local tradables through improving market access. Thus, one contribution of our paper is to generalize the predictions from the two literatures that have focused on either the costs or the benefits from booming sectors.

Our paper is also related to the quantitative international trade literature, which has developed tractable models of cross-country trade with various dimensions of heterogeneity (Anderson, 1979; Eaton and Kortum, 2002; Melitz, 2003). These seminal models characterize trade and the distribution of economic activity as a function of exogenous trade costs. A standard prediction of these models is that the relationship between trade flows and costs follows a gravity equation, which has been documented as one of the strongest empirical regularities in the data (Head and Mayer, 2014). We complement this literature by developing a framework in which trade costs are *endogenous*, in a way that is both tractable and preserves the gravity structure of trade flows. This relates our paper to Fajgelbaum and Schaal (2020) and Santamaría (2022), who consider endogenous road construction in multi-location models of economic geography, as well as Brancaccio, Kalouptsidi, and Papageorgiou (2020), who endogenize trade costs in the non-containerized shipping sector. Unlike these papers, we focus on port development as a source of endogenous shipping costs, and solve for the decentralized equilibrium as opposed to the optimal allocation to quantify the effects of port development on trade, the distribution of population, and welfare.

Finally, our paper is related to a large literature studying the effects of transport infrastructure improvements.[3] Within this literature, Brinkman and Lin (2022) is the only other paper we are aware of that shows empirical evidence for the cost side of infrastructure development, focusing on disamenities associated with freeway construction in mid-20th century U.S. cities. Our paper also relates to the growing empirical literature studying the effects of containerization (Hummels, 2007; Bernhofen, El-Sahli, and Kneller, 2016; Gomtsyan, 2016; Coşar and Demir, 2018; Holmes and Singer, 2018; Altomonte et al., 2018; Brooks et al., 2021; Bridgman, 2021) or the role of

[2]Relatedly, Falvey (1976) discusses how the transportation sector can draw away resources from tradables in particular.

[3]Redding and Turner (2015) provides an overview of recent developments in this literature. Ducruet and Notteboom (2023) reviews the existing port geography literature.

container shipping networks in world trade (Wong, 2022; Heiland, Moxnes, Ulltveit-Moe, and Zi, 2022; Ganapati, Wong, and Ziv, 2022; Koenig, Pigné, Poncet, Sanch-Maritan, and Duvallet, 2023). Most closely related in this vein is Brooks et al. (2021), who study the reduced-form effects of containerization on local economic outcomes across U.S. counties. Our main contribution to this literature is twofold. First, our paper is the first to highlight that containerization leads to sizeable local and global costs.[4] Second, to the best of our knowledge, ours is the first paper seeking to quantify the aggregate effects of containerization on global trade and welfare through the lens of a general equilibrium economic geography model.

The paper is structured as follows. In the next section, we describe the transshipment cost reductions caused by containerization. Section 2 discusses the main data sources used in the analysis. Section 3 presents three stylized facts about the local effects of containerization, while Section 4 introduces the model. Section 5 estimates a model guided empirical specification that disentangles the local costs and benefits of port development. Section 6 measures the aggregate effects of containerization, and Section 7 illustrates the effects of targeted port development policies similar to the Maritime Silk Road. Finally, Section 8 concludes.

1 Background: containerization reduced transshipment costs

As late as the mid-1950s, transshipment at seaports was a costly and slow procedure as it entailed handling cargo item-by-item – a process called breakbulk shipping (Krugman, 2011). Cargo came in many different sizes and needed to be handled individually, despite the widespread use of machinery introduced pre-containerization (see Panel A of Figure E.1). The San Francisco Port Commission (1971) estimated that it took 7 to 10 days to merely discharge cargo from a ship. According to Bernhofen et al. (2016), two-thirds of a ship's time was spent in port. This led to high costs as the capital utilization of ships was low, and the cost of capital tied up in inventory was high.[5]

U.S. shippers first started placing cargo into containers in the late 1950s. Containerized shipping was initially introduced on domestic routes between U.S. ports, but the technology was rapidly adopted and standardized worldwide over the next two decades (Rua, 2014). Containerized port technology can be seen in its mature form at the Port of Seattle in 1969 in Panel B of Figure E.1 (a mere 10 to 15 years after the photos shown in Panel A were taken). Cargo, packed in standardized containers, is loaded onto and off ships using large, purpose-built cranes situated on the wharf. Large, open areas beside the wharf are used to line up containers.

[4]Our result that land-abundant cities see faster shipping growth is consistent with Brooks et al. (2021) who find that containerization led to faster population growth in U.S. counties with initially low land rents.

[5]Industry experts estimated that the handling of cargo at the port accounted for a major share of freight costs (Levinson, 2010). As an example, transshipment costs were estimated to account for 49% of the total transport cost on one route from the U.S. to Europe (Eyre, 1964).

Containerization substantially reduced transshipment costs for a number of reasons. First, as containers could be handled in a uniform way, loading and unloading times were vastly reduced. The San Francisco Port Commission (1971) estimated that a container ship could be unloaded and loaded in 48 hours or less, a tenth of the previous time spent in port. Similarly, using detailed data on vessel turnaround times for an anonymized port, Kahveci (1999) estimates that the average time ships spent in port fell from 8 days to 11 hours as a result of containerization, a reduction of 94%. Second, the reduction in turnaround time justified investment in much larger vessels (Gilman, 1983). The average size of newly-built container ships increased by 402% between 1960 and 1990.[6] Larger ship sizes made it possible to realize even larger cost reductions through increasing returns to scale in shipping and port handling. Rodrigue (2016, p. 118) estimates that moving from a 2,500 TEU capacity vessel to one with 5,000 TEU reduced costs per container by 50%.

2 Data

Our analysis builds on a decadal city-level dataset of shipping flows, population, and other economic outcomes for the period 1950-1990. We complement this with GIS data that allows us to calculate geographic characteristics of the cities and ports. We review the main variables used in the analysis below and report summary statistics in Table D.1. Detailed documentation including sources and description of data construction for all the data used in the paper can be found in the Data Appendix (Appendix C).

Shipping flows (Appendix C.1). Crucial to our analysis is a unique dataset of worldwide bilateral ship movements at the port level for the period 1950-1990 from Ducruet, Cuyala, and Hosni (2018). An observation is a ship moving from one port to another at a particular point in time.[7] One week samples from the first week of May for each year were extracted from the *Lloyd's Shipping Index*, a unique source that provides a daily list of merchant vessels and their latest inter-port movements.

These data provide us with rich variation to study the geography of sea-borne trade through the second half of the 20th century. They cover both domestic and international shipping. Moreover, the data cover a long time period spanning the containerization revolution. We are thus able to compare the effects of port activity on cities both before and after the arrival of the new technology. We know of no other data source that has a similar coverage across time and space, especially at such a detailed level of disaggregation. An important limitation, however, is that we do not observe either the value or the volume of shipment but only bilateral ship movements. From these ship

[6] These calculations are based on data from the *Miramar Ship Index* (Haworth, 2020). See Appendix C.11 for details.

[7] As such, it is similar to contemporary satellite AIS (Automatic Identification System) data that tracks the precise movements of vessels around the globe. Such AIS data are used in Brancaccio et al. (2020) and Heiland et al. (2022).

movements, we sum the total number of ships passing through each port, which we call *shipping flows*.

City population. As we are interested in the economic effects of containerization, we use data on city population worldwide for locations with more than 100,000 inhabitants from *Villes Géopolis* (Moriconi-Ebrard, 1994) for each decade between 1950-1990 (Geopolis cities, henceforth). The advantage of these data relative to sources such as the more frequently used *UN World Cities* dataset is that a consistent and systematic effort was made to obtain populations for the urban agglomeration of cities (that is, the number of inhabitants living in a city's contiguous built-up area) as opposed to the administrative boundaries that are often reported in country-specific sources. This definition of the city ensures that the port lies within the city boundaries even if it is outside the administrative boundaries of the city. For example, New York (New York) and Elizabeth (New Jersey), which includes the port of Elizabeth, form one 'city' according to this definition. We observe population for cities that reached 100,000 inhabitants in any year throughout this period. For most of these cities, we observe population even when the city had fewer than 100,000 inhabitants, potentially leading to sampling bias. To address this, we will show that our results are robust to using the subset of cities that had already attained 100,000 inhabitants in the first sample year, 1950.

Ports were hand-matched from the shipping data to cities based on whether the port was located within the urban agglomeration of a city in the Geopolis dataset, allowing for multiple ports to be assigned to one city (Ducruet et al., 2018). We define port cities in a time-invariant manner; a port city with positive shipping flows in at least one year will be classified as a port city for all years. Of the 2,636 cities in the Geopolis dataset, 553 have at least one port. We label these as *port cities*. The quantitative estimation covers the full set of 2,636 Geopolis cities (port and non-port cities).

Underwater elevation levels (Appendix C.2). We use gridded bathymetric data on underwater elevation levels at a detailed spatial resolution (30 arc seconds, or about 1 kilometer at the equator) from the *General Bathymetric Chart of the Oceans (GEBCO)* to measure sea depth around port cities.

Saiz land scarcity measure (Appendix C.3). To measure cities' land scarcity, we follow the methodology in Saiz (2010), using GIS data that have global coverage: We take a 50 kilometer radius around the centroid of the city, and count all sea cells, all internal water bodies and wetland areas, as well as all cells with a gradient above 15%. These cells, as a share of the total cells, can be used as a proxy for a city's land scarcity, as they cannot be built on.[8]

City-level GDP per capita (Appendix C.4). Data on city-level income levels are needed for the

[8]Saiz (2010) argues that this measure (or rather, 1 minus our measure) captures land supply well, as it is positively correlated with rents in his sample.

quantitative estimation only. We are not aware of readily available sources of GDP per capita data for cities worldwide. For this reason, we estimate GDP per capita for the last year in our sample (1990) for the full sample of 2,636 worldwide cities in the following way. First, we use estimates of city GDP from the *Canback Global Income Distribution Database* for a subset of our sample (898 cities) for which data are reported for 1990. We extrapolate GDP per capita for the full sample of cities using the linear fit of the GDP per capita data on nightlight luminosity and country fixed effects, building on a growing body of evidence suggesting that income can be reasonably approximated using nightlight luminosity data (Donaldson and Storeygard, 2016).

Google Earth port area (Appendix C.7). To measure the land area of ports, we hand-coded polygons from *Google Earth* that we identified as containing port activities for a random set of 236 port cities in our dataset.

3 Stylized facts

In this section, we document three stylized facts about the local effects of containerization on port cities. First, we show that containerization led to an increase in shipping flows for port cities that were exogenously more suited to containerization, because they have deeper ports. Second, we show that there are heterogeneous effects: The increase in shipping was larger for port cities in which land is less scarce. Third, we show that the increase in shipping did not translate into population growth. Together, these stylized facts suggest that containerization entailed both costs and benefits for host cities.

3.1 Stylized fact 1: Containerization led to shipping growth in deeper port cities

Section 1 discussed the fact that containerization led to larger ship sizes. This, in turn, required greater depth at the port. Following the previous literature, we think of *naturally endowed* depth as an exogenous cost-shifter that makes it cheaper for a port to reach a desired depth through costly dredging (Brooks et al., 2021; Altomonte et al., 2018). The empirical challenge is that *observed* port depth is a combination of naturally endowed depth and depth attained by dredging. Our solution to this relies on using contemporary granular data on underwater elevation levels around the port to isolate the naturally endowed component of depth. In particular, we take all sea cells within buffer rings around the geocode of the port and sum the number of cells that are 'very deep,' which we define as depth greater than 30 feet following Brooks et al. (2021). These authors argue that given vessel sizes in the 1950s (pre-containerization), depth beyond 30 feet conferred no advantage to the port. Below, we will test how reasonable this assumption is by examining pre-trends in shipping. Our baseline measure of port suitability is thus the log of the sum of 'very deep' cells in a buffer ring 3-5 km around the port.[9] We examined the effect of depth measured

[9]There are zeros in the data, that is, there are ports with no cells deeper than 30 feet in the 3-5 km buffer around the port. For this reason, in practice, we use $\ln(1 + \sum_i \mathbb{1}(depth_i \geq 30ft))$, where i denotes a cell.

at various buffers and confirmed that the effects are similar in nearby rings, suggesting that the variation we use from the 3-5 km buffer is a representative measure of depth at the port.

The key assumption behind our ability to isolate naturally endowed depth (from depth attained by dredging) is that when ports need to invest in costly dredging, they typically do not dredge entire areas in our buffers, but narrow channels that ships use to navigate to the port. By calculating depth over many sea cells, the vast majority of depth measurements for each port should reflect naturally endowed depth. We test this assumption in the following way. For 100 random ports in our sample, we obtained access to nautical maps from *marinetraffic.com*. These clearly demarcate the dredged channels that ships use to navigate to the port.[10] We then construct a binary variable, *'Dredging'*, that takes the value of 1 if a port has a dredged channel in the 3-5 km buffer ring. Table D.2 shows the association between this measure and the depth measure. The unconditional association (column 1) is *negative* and statistically significant. That is, ports that we measure to be shallow are more likely to have a dredged channel. This is what we would expect to find if our measure captured naturally endowed depth.[11]

The following flexible specification allows us to estimate the causal effect of containerization on shipping, driven by exogenously endowed port depth.

$$\begin{aligned}
\ln(Ship_{it}) \quad = \quad & \sum_{j=1960}^{1990} \beta_j * Depth_i * \mathbb{1}(Year = j) + \\
& \sum_{j=1960}^{1990} \phi_j * \ln(Pop_{i,1950}) * \mathbb{1}(Year = j) + \alpha_i + \delta_t + \epsilon_{it}
\end{aligned} \tag{1}$$

The outcome variable of interest, $\ln(Ship_{it})$, is the log of shipping flows observed in city i at time t. We need to take a stand on the treatment of zeros in the shipping data. In the baseline measure, we annualize the weekly counts of ships from the raw data by multiplying the one-week sample of shipping flows we observe by 52. This is primarily so that our results are comparable to regressions we run using model-simulated data in the quantification exercise in Section 6. Finally, we replace the zeros in the data with ones and take the natural logarithm of this adjusted annualized count.[12]

$Depth_i$ is the cross-sectional measure of port suitability defined in the previous subsection. We

[10] We provide more details on this exercise in Appendix C.12.

[11] Adding continent or coastline fixed effects (columns 2 and 3, respectively) reduces the size of the negative coefficient and we lose statistical significance in column (3), but the estimated coefficients remain negative.

[12] In robustness checks discussed below, we show that all of the results presented in this section are robust to other standard ways of dealing with the zeros. In these, we do not annualize the data in order to verify that this transformation does not drive the results.

interact this measure with binary indicators for the decades 1960 – 1990 to estimate the time path of how depth affected shipping flows. Since containerization spread out globally towards the end of the 1960s, when international standards for the size of containers were introduced, we would expect depth to positively affect shipping only after 1970. We include the full set of city and year fixed-effects (denoted α_i and δ_t, respectively), and also allow for the initial population in 1950 to have a time-varying effect on shipping. The latter ensures that we do not mistake population convergence patterns, i.e., initially smaller cities experiencing stronger growth, as the effect of containerization.[13] We cluster standard errors at the city level in the baseline to account for the serial correlation of shocks. We also report Conley standard errors (in curly brackets).[14] Each β_j in this specification estimates the increase in shipping caused by having a deeper port in a given year relative to 1950.

Table 1 contains the estimated coefficients. Column (1) presents coefficients for the baseline specification. Consistent with containerization technology being rolled out in the early 1960s across US ports and worldwide later in the decade, we see that deeper ports experienced differential growth in shipping flows only from 1970 onwards, but not in the decade between 1950 and 1960. The effect of depth on shipping is much larger and significantly different from zero for the interaction of depth and each year indicator including and after 1970.

A causal interpretation of the estimated effect of depth relies on the identifying assumption that the time-varying effect of depth is uncorrelated with the error term. The timing of when depth started to matter and the lack of pre-trends provide evidence that this assumption is plausible. Next, we turn to further testing this result with more demanding specifications. One concern is that many determinants of depth may be spatially correlated and if true, the estimates could be hard to disentangle from broader regional trends. To this end, column (2) adds the full set of 'coastline' by year-fixed effects to examine the extent to which our identifying variation relies on cross-regional variation.[15] Note that this set of fixed effects subsumes continent by year fixed effects.

Since geological features are correlated across space, deeper ports may be correlated with a more rugged terrain in the city. As we will argue below, land-scarce cities are less able to take up the opportunities of containerization, so we would expect this correlation to bias our estimates

[13]If we did not control for initial population, the estimated effects would be slightly larger, with no statistically significant pre-trends, but we prefer to be conservative and not attribute this part of the growth to containerization.

[14]As these are typically very close to the clustered standard errors, we only report them for the main results for easier readability of the tables. We allow for spatial correlation at distances up to 1,000 km and set the spatial decay function to be linear.

[15]We define coastlines in the following way. We assign each port to its nearest ocean (e.g., 'Pacific Ocean') or body of water (e.g., 'Great Lakes') and further disaggregate oceans by continent. This yields 22 coastlines worldwide. Examples are 'Mediterranean – Europe' and 'North America – Atlantic.'

towards zero. This is confirmed by column (3), where we add the Saiz land scarcity measure interacted with year indicators to allow for differential trends across more and less land-scarce port cities. As expected, the estimated effect of port depth on shipping increases. Column (4) adds country GDP per capita (measured in 1960) interacted with year indicators to control for potentially differential growth trends across initially rich and poor countries.[16]

Overall, there is a consistent absence of pre-trends, and a consistent effect of depth on shipping in the years 1970 and after. Based on these results, we introduce a 'containerization' treatment indicator that turns on in years including and after 1970. This yields a single coefficient that estimates the differential effect of depth on shipping after the onset of containerization. Column (5) shows the results. Cities endowed with more depth, and hence more suitable to containerized technologies witnessed disproportionate increases in their shipping flows after containerization. Panel A in Figure E.2 shows that the coefficient of interest remains fairly stable as we drop continents one at a time, underscoring that no single region appears to be driving the results. The coefficient becomes somewhat smaller when we drop North America, which is in line with the United States being the birthplace and an early adopter of containerization.

Panel A in Table D.3 contains further robustness checks. First, we test robustness to different data construction choices. In particular, we examine different ways of treating zero shipping values, different ways of defining the depth measure for the handful of ports that are located far inland from the coastline and restricting the sample to the subset of cities that had already attained 100,000 inhabitants by 1950 to examine sample selection bias. The coefficient of interest remains similar in magnitude and highly significant across all these checks. We now turn to examining whether this containerization-induced shipping boom was heterogeneous across port cities.

3.2 Stylized fact 2: Ports expanded more in response to containerization where land was less scarce

Modern container ports require vast amounts of land. Faster turnaround times can only be achieved by building much larger terminals. Rodrigue (2016, p. 118) names site constraints, and in particular, the large consumption of terminal space as the primary challenge associated with containerization. In this section, we first document the increased land-intensity of containerization in historical and contemporaneous data. Next, we examine how the increased land intensity affected *where* port development took place.

The increased land-intensity of containerized ports. Historical case study evidence from a number of ports shows that successful containerization required substantial geographic expansion of the port. In a 1971 report, alarm bells were rung about the inadequacy of San Francisco's finger piers

[16] We use the 1960 (pre-containerization) measure of country GDP per capita as this is observed for a larger set of countries than for 1950.

to accommodate new types of cargo handling; “No pier facilities in the Bay Area today are capable of handling the new space requirements on this scale of new and larger container ships. (...) thus more berthing and backup area is needed” (1971, p. 13). Ports such as the one in San Francisco that were adjacent to a densely built up city struggled (and often ultimately failed) to find the necessary space for container port development (Corbett, 2010, p. 164). In contrast, at ports where containerization succeeded, the port expanded substantially. Using detailed, annual engineering maps and cargo throughput for the Port of Seattle, we find that the area of the port increased fourfold, while the land intensity of the port (i.e., the area of the port relative to throughput) almost doubled between 1961-1973, the period when the port containerized (see Appendix A for further historical evidence and Appendix C.6 for a discussion of Seattle’s containerization, respectively).

The land intensity of containerized terminals is also evident in contemporaneous data. We exploit the availability of high resolution remote sensing data that makes it possible to measure the area of ports in our sample today. We match this with data on the cargo composition handled by each port using data from *Le Journal de la Marine Marchande (JMM)* for 2008-09. Column (2) in Table D.4 shows that, controlling for the total volume of traffic, ports that handle more containerized cargo are typically larger. In columns (3)-(6), we control for additional potential confounders. The coefficient remains significant and very similar in magnitude when we add controls for the volume of non-liquid and solid bulk handled by the port (e.g., oil, coal or grain, the handling of which may be technologically very different), the GDP per capita of the country (to get at differences in the extent of automation), or even country fixed effects.

To get a better sense of magnitudes, in column (7) we use the *share* of containerized cargo (while continuing to control for the total volume of cargo). The binscatter (plotted in Figure E.3) visualizes the positive relationship between the area occupied by the port and the share of containerized cargo. The coefficient of interest is large and statistically significant. Based on this specification, moving from a fully non-containerized port to a fully containerized port is associated with a 75% increase in port area (i.e., $exp(0.5624) - 1$), holding the volume of traffic fixed.

Port development took place where land was less scarce. Across cities, we can examine whether shipping increased more in cities where land was less scarce by allowing for heterogeneous effects with respect to land scarcity in regression equation (1):

$$\begin{aligned}\ln(Ship_{it}) \quad = \quad & \beta * Depth_i * \mathbb{1}(Year \geq 1970) + \gamma * Depth_i * LandScarcity_i * \mathbb{1}(Year \geq 1970) \\ & + \eta * LandScarcity_i * \mathbb{1}(Year \geq 1970) + \sum_{j=1960}^{1990} \phi_j * \ln(Pop_{i,1950}) * \mathbb{1}(Year = j) \\ & + \alpha_i + \delta_t + \epsilon_{it} \qquad (2)\end{aligned}$$

where $LandScarcity_i$ measures the share of land in city i that cannot be built on, as defined in Section 2, and all other variables are as defined above. We have defined the measure such that higher values correspond to a city with more land scarcity. The coefficient of interest is γ—the interaction between our depth measure and the land scarcity measure (interacted with the 'containerization' treatment variable that turns on in 1970). Note that this is a fully saturated specification. We allow both depth and the land scarcity measure to have their own time trend break in 1970.

We plot the marginal effect of depth at different values of the land scarcity measure in Figure 1a (the corresponding estimates are presented in Table D.5). Consistent with an important role for land scarcity in determining the location of port development, the coefficient of interest, γ, is negative, large and statistically different from zero (coefficient -0.707, s.e. 0.323). Cities with exogenously deeper ports witnessed increased shipping flows after 1970, but disproportionately more so in cities where land was less scarce.

Our land scarcity measure is constructed using contemporary GIS data, which captures natural geography in combination with investments in reclaiming land from the sea. To investigate the extent to which reclamation may introduce systematic measurement error, we use data from Martín-Antón, Negro, López-Gutiérrez, and Esteban (2016) on coastal land reclamation conducted for any potential purpose.[17] Table D.6 shows that there is somewhat more land reclamation in cities we measure to be land-scarce. This is what we would expect if the land scarcity measure was mostly capturing natural geography. The reason for this seems to be that while land reclamation is fairly common (76 out of 553 ports report *some* land reclamation), it is typically small relative to the area over which the land scarcity measure is constructed.[18] One may also expect that land reclamation is easier in shallower ports. However, the estimated relationship between the depth measure and the binary indicator of land reclamation is small and never statistically significant.

Panel B in Figure E.2 explores the heterogeneity of the result by dropping continents one at a time. The effect is consistently negative. We perform the same set of robustness checks for this result as for the previous one (see Tables D.3, panel B, and D.5). The results are largely robust to these specifications, as our coefficient of interest, γ, remains negative and economically large throughout all these checks. Finally, we note that over time (1953 – 2017), ports systematically moved *within city* to the outskirts, where land is typically less scarce (Figure E.4).

3.3 Stylized fact 3: The increase in shipping did not translate into population growth

To document the long-run effect of containerization-induced port development on population, we estimate the following long-differenced specification;

[17] See Appendix C.8 for a discussion of the data.

[18] The median size of reclaimed area in the sample for the non-zero observations is 13 square kilometers, which pales in comparison to the 7850 square kilometers covered in the land scarcity measure.

$$\Delta \ln(Pop_i) = \beta * \Delta \ln(Ship_i) + \phi * \ln(Pop_{i,1950}) + \epsilon_i \quad (3)$$

where $\Delta \ln(Pop_i)$ and $\Delta \ln(Ship_i)$ are the change in the natural logarithm of population and shipping flows between 1950 and 1990, respectively. The identification challenge is that the shipping flows of a city are endogenous. Our main worry is reverse causality: fast growing cities will witness increases in their shipping flows. Our solution is to isolate exogenous variation in shipping using a city's suitability for containerization based on its natural depth. We control for initial population levels to account for population convergence.

Table 2 contains the baseline regression results. Both the estimated OLS and 2SLS coefficients on shipping are small and statistically indistinguishable from zero (OLS coefficient 0.013, s.e. 0.009; 2SLS coefficient 0.006, s.e. 0.073). To assess magnitudes, we report the standardized 'beta' coefficients for our effects of interest in italics underneath the estimated regression coefficients. A one standard deviation increase in the growth of shipping flows between 1950 and 1990 leads to a 0.02 standard deviation increase in population growth over the same time horizon based on the 2SLS estimate. Columns (3) and (4) show the first stage and reduced form, respectively. These help illuminate what drives the small and insignificant effect. While the first stage coefficient is highly significant and the Kleibergen-Paap F-statistic is reasonable (9.98), there is no reduced form relationship between depth and population (the reduced form coefficient is 0.002, s.e. 0.020).

Table D.7 shows the panel specification allowing us to utilize the full decadal variation in the data.[19] Two important points emerge. First, the results are very similar to the long-differenced specification. The 2SLS coefficient remains small in magnitude and statistically indistinguishable from zero. The first stage is strong (the Kleibergen-Paap F-statistic is 21.13), and the reduced form is small and statistically insignificant. Second, column (5) shows the full time path of effects for the reduced form. These make clear that the statistically insignificant coefficient in the 2SLS estimate does not stem from the fact that population is sluggish to adjust. The time path of the coefficients shows no discernible trend, and there is no clear difference in population growth post-containerization for deeper ports. All of the coefficients are estimated to be very close to zero (the one 'furthest' away from zero is 0.007), the coefficients are never close to statistical significance, and in two of the five decades, the estimated effect is negative, suggesting that, if anything, deeper ports were growing at a slower rate than shallower ones some of the time.

We subject the 2SLS panel specification to the same set of robustness checks conducted above: we include coastline by year fixed effects, as well as control for the time varying-effects of the land scarcity measure and GDP per capita (Table D.8). The coefficient is consistently small and

[19] The specification is $\ln(Pop_{it}) = \beta * \ln(Ship_{it}) + \sum_{j=1960}^{1990} \phi_j * \ln(Pop_{i,1950}) * \mathbb{1}(Year = j) + \alpha_i + \delta_t + \epsilon_{it}$, where $\ln(Pop_{it})$ is the natural logarithm of population in city i at time t, and all other variables are as previously defined.

indistinguishable from zero. No single continent drives this result (Figure E.2).[20] Table D.3, panel C, shows that the results are robust to the same set of additional robustness checks to data choices performed in the other panels.

In summary, these results show that we cannot reject that the effect of increased port activity on population was zero. Given that increased trade through a city tends to increase population through the standard market access effect (Donaldson and Hornbeck, 2016; Redding and Turner, 2015), this finding suggests a role for countervailing forces. The large standard errors typical of 2SLS estimation do not allow for a definitive answer based on this evidence alone. However, guided by the model, we will be able to formally test for the presence of negative as well as positive forces in Section 5.

4 A model of cities and endogenous port development

The three stylized facts documented in Section 3 suggest that while containerization had a positive effect on shipping, and more so in cities that were less land-scarce, port development also entailed local costs for the host cities, as the increase in shipping did not translate into population growth. In this section, we present a flexible general equilibrium model that is consistent with these three stylized facts and allows us to estimate the aggregate and distributional effects of port development. The model captures the standard positive effects from market access, but also allows for two types of negative effects: the increased land use and the negative amenity externalities associated with port development.

4.1 Setup

The world consists of $S > 0$ cities, indexed by r or s. An exogenously given subset of cities are port cities, while the rest are non-port cities. We make the Armington assumption that each city produces one variety of a differentiated final good that we also index by r or s (Anderson, 1979). Each city belongs to one country, and each country is inhabited by an exogenous mass of workers who choose the city in which they want to live. We do not allow for mobility across countries but allow for mobility across cities within a country, subject to frictions.

[20] This exercise also helps reconcile the results presented here with Brooks et al. (2021) who find a positive effect of containerization on county population growth in the United States. A direct comparison is not possible as our sample only contains 40 U.S. cities and the 2SLS estimate on this subsample yields a Kleibergen-Paap F-statistic below 1. However, dropping North America leads to a negative (though statistically insignificant) point estimate (Figure E.2), suggesting that on average North American cities may have had a larger than average population response to containerization.

4.1.1 Workers

Each worker owns one unit of labor that she supplies in her city of residence. The utility of a worker j who chooses to live in city r is given by

$$u_j(r) = \left[\sum_{s=1}^{S} q_j(r,s)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} a(r)\, b_j(r) \tag{4}$$

where $q_j(r,s)$ is the worker's consumption of the good made in city s, $a(r)$ is the level of amenities in city r, and $b_j(r)$ is an idiosyncratic city taste shifter. $\sigma > 1$ is the elasticity of substitution across goods.

The dispersion of $b_j(r)$ represents the severity of cross-city mobility frictions that workers face, similar to Kennan and Walker (2011) and Monte, Redding, and Rossi-Hansberg (2018). For tractability, we assume that $b_j(r)$ is drawn from a Fréchet distribution with shape parameter $1/\eta$ and a scale parameter normalized to one. Hence, a larger value of η corresponds to more severe frictions to mobility.

We also capture the fact that port activity might induce disamenities such as noise and pollution. In particular, we assume

$$a(r) = \bar{a}(r)\left[1 + Shipping(r)\right]^{-\rho} \tag{5}$$

where $\bar{a}(r)$ is the city's fundamental, exogenous amenity level, $Shipping(r)$ is the total amount of shipping flowing through the port of city r, and $\rho > 0$. In non-port cities, by definition, $Shipping(r) = 0$, implying $a(r) = \bar{a}(r)$. In port cities, fundamental amenities $\bar{a}(r)$ are lowered by the term $[1 + Shipping(r)]^{-\rho}$, implying that a larger volume of shipping is associated with more disamenities. The extent to which this is the case is disciplined by the value of parameter ρ.

4.1.2 Landlords

Each city r is also inhabited by a positive mass of immobile landlords who own the exogenously given stock of land available in the city.[21] We normalize the stock of land available in each city to one.[22] Landlords have the same preferences over goods as workers. They do not work but finance

[21]The assumption about the elasticity of land supply merits further discussion. A perfectly elastic land supply would not yield a land use cost of port development as cities would respond to containerization by expanding their stock of land. As we find empirical evidence in support of sizeable local costs from containerization (Section 5) that cannot be explained by disamenities alone (Section 6.4), we need to move away from the case of perfectly elastic land supply. To retain the tractability of the model, we assume that land supply is perfectly inelastic and leave the case of imperfectly elastic supply for future research.

[22]We could allow the stock of available land to vary across cities. This more general setup is isomorphic to our current model, except that, instead of productivity in the city-specific good sector, a combination of the stock of land and productivity enters the model's equilibrium conditions. In other words, the city productivity levels we identify from our current model reflect not only productivity per se, but also the stock of available land. This fact, however, does not affect our quantitative results as we keep productivity levels

their consumption from the revenues they collect from their stock of land.

Each landlord is small relative to the total mass of landlords in the city and hence thinks that she cannot influence prices. Yet the mass of landlords is small enough that the population of each city can be approximated well with the mass of workers who choose to reside in the city.

In non-port cities, landlords rent out their land to firms that produce the city-specific good. In port cities, landlords allocate their land between what they rent out to firms for production and what they use for transshipment services at the port. The more land they use for transshipment services, the more the cost of transshipping a unit of a good decreases, as we explain in detail in Section 4.1.4. The landlord can charge a price for the transshipment service she provides. Competition among port city landlords drives down this price to marginal cost. Hence, profits from transshipment services are zero in equilibrium.[23]

4.1.3 Production

Firms can freely enter the production of the city-specific good. Hence, they take all prices as given and make zero profits. Production requires labor and land. The representative firm operating in city r faces the production function

$$q(r) = \tilde{A}(r)\, n(r)^{\gamma} (1 - F(r))^{1-\gamma}$$

where $q(r)$ denotes the firm's output, $\tilde{A}(r)$ is total factor productivity in the city, $n(r)$ is the amount of labor employed by the firm, and $F(r)$ is the share of land that landlords in the city allocate to transshipment services (thus, $F(r) = 0$ in non-port cities). Hence, $1 - F(r)$ is the remainder of land that landlords rent out to firms for production, and γ and $1 - \gamma$ correspond to the expenditure shares on labor and land, respectively.

We incorporate agglomeration economies by allowing total factor productivity to depend on the population of the city, $N(r)$:

$$\tilde{A}(r) = A(r)\, N(r)^{\alpha}$$

where $A(r)$ is the exogenous fundamental productivity of the city, and $\alpha \in [0, 1-\gamma]$ is a parameter that captures the strength of agglomeration economies.[24] The representative firm does not internalize the effect that its employment decision has on local population. Hence, it takes $N(r)$

---

fixed in our model simulations.

[23] In Section 6, we show that the effects of containerization remain similar in an alternative framework in which landlords have market power and thus can make profits. We provide a detailed description of this alternative framework in Appendix B.7.

[24] We make the assumption $\alpha \leq 1 - \gamma$ to guarantee that agglomeration forces are not overwhelmingly strong in the model. Estimates of the land share, $1 - \gamma$, tend to be substantially above estimates of agglomeration externalities α. In particular, our calibration involves setting α to 0.06 (a standard value used in the literature) and $1 - \gamma$ to 0.16 based on Desmet and Rappaport (2017).

as given.

4.1.4 Shipping and port development

Firms in city r can ship their product to any destination $s \in S$. Shipping is, however, subject to iceberg costs: if a firm i from city r wants to ship its product over a route $\bar{\rho}$ that connects r with s, then it needs to ship $T(\bar{\rho}, i)$ units of the product such that one unit arrives at s. Shipping costs consist of a component common across firms $\bar{T}(\bar{\rho})$, as well as a firm-specific idiosyncratic component $\epsilon(\bar{\rho}, i)$ that is distributed i.i.d. across firms and shipping routes:[25]

$$T(\bar{\rho}, i) = \bar{T}(\bar{\rho})\, \epsilon(\bar{\rho}, i)$$

For tractability, we assume that $\epsilon(\bar{\rho}, i)$ is drawn from a Weibull distribution with shape parameter θ and a scale parameter normalized to one. Firms only learn the realizations of their idiosyncratic cost shifters after making their production decisions. Therefore, they make these decisions based on the expected value of shipping costs,

$$\mathbf{E}\left[T(\bar{\rho}, i)\right] = \bar{T}(\bar{\rho})\, \mathbf{E}\left[\epsilon(\bar{\rho}, i)\right] = \bar{T}(\bar{\rho})\, \Gamma\left(\frac{\theta+1}{\theta}\right).$$

After learning $\epsilon(\bar{\rho}, i)$, they choose the route that minimizes their total shipping costs.

Certain shipping routes involve land shipping only (*land-only*), while others involve a combination of land and sea shipping through a set of ports (*land-and-sea*). Land-only shipping is only available between cities that are directly connected by land. The common cost of land-only shipping between cities r and s is an increasing function of the minimum overland distance between the two cities, $d(r, s)$:

$$\bar{T}(\bar{\rho}) = 1 + \phi_{\varsigma}(d(r, s))$$

The cost of land-and-sea shipping depends on the set of ports en route. In particular, the common cost of shipping from r to s through port cities $p_0, ..., p_M$ takes the form

$$\bar{T}(\bar{\rho}) = \left[1 + \phi_{\varsigma}(d(r, p_0))\right]\left[1 + \phi_{\varsigma}(d(p_M, s))\right] \prod_{m=0}^{M-1}\left[1 + \phi_{\tau}(d(p_m, p_{m+1}))\right] \prod_{m=0}^{M}\left[1 + O(p_m)\right]$$

where $\phi_{\varsigma}(d(r, p_0))$ corresponds to the overland shipping cost between the origin and the first port en route p_0, and $\phi_{\varsigma}(d(p_M, s))$ corresponds to the overland shipping cost between the last port en route p_M and the destination. $\phi_{\tau}(d(p_m, p_{m+1}))$ denotes the sea shipping cost between ports p_m

[25] The assumption of idiosyncratic shipping cost shifters follows Allen and Atkin (2022) and Allen and Arkolakis (2019), and allows us to tractably characterize shipping flows with a large number of cities. In the alternative case with no idiosyncratic shifters, applied in Allen and Arkolakis (2014) and Nagy (2022), finding optimal shipping flows is computationally more demanding.

and p_{m+1}, a function of the minimum sea distance between the two ports, $d(p_m, p_{m+1})$. Finally, $O(p_m)$ denotes the price that the firm needs to pay for transshipment services in port city p_m.[26]

Transshipment costs are central to our analysis as these are the costs that port city landlords can lower through *port development*, that is, through allocating more land to the port. In particular, we assume that the landlord's cost of handling one unit of a good at port p_m equals

$$[\nu(p_m) + \psi(F(p_m))]\, Shipping(p_m)^{\lambda}$$

where $\nu(p_m)$ is an exogenous cost shifter capturing the fundamental efficiency of port p_m, $\psi(F(p_m))$ is a non-negative, strictly decreasing and strictly convex function of $F(p_m)$, the share of land allocated to the port, and $Shipping(p_m)^{\lambda}$ captures congestion externalities arising from the fact that handling one unit of cargo becomes more costly as the total amount of shipping, $Shipping(p_m)$, increases for a given port size.[27] As each port city landlord is atomistic, she takes the price of transshipment services $O(p_m)$ and the total port-level shipping $Shipping(p_m)$ as given when choosing $F(p_m)$. Moreover, perfect competition among port city landlords ensures that the price of transshipment services is driven down to marginal cost and therefore

$$O(p_m) = [\nu(p_m) + \psi(F(p_m))]\, Shipping(p_m)^{\lambda} \tag{6}$$

in equilibrium.

One concern is that, according to our formulation, land is required for transshipment services while labor is not. In reality, ports employ labor. To address this concern, Appendix B.6 presents an extension of our model in which a combination of land and labor must be employed in transshipment. This appendix also shows that the model with transshipment labor, although more complex in its structure, delivers qualitative predictions that are extremely similar to the predictions of our baseline model.

4.1.5 Equilibrium

In equilibrium, workers choose their consumption of goods and residence to maximize their utility, taking prices and wages as given. Landlords choose their consumption and land use to maximize their utility, taking prices, land rents and shipping flows as given. Firms choose their production of goods, employment and land use to maximize their profits, taking prices, land rents and wages as

[26]Note that this formulation does not allow for land shipping between two subsequent ports along the route. In practice, this is extremely unlikely to arise as land shipping is substantially more expensive than sea shipping.

[27]To be precise, $Shipping(p_m)$ is defined as the dollar amount of shipping flowing through port p_m, excluding the price of transshipment services at p_m. We exclude the price of transshipment services from the definition of $Shipping(p_m)$ as it simplifies the procedure of taking the model to the data.

given. Competition drives profits from production and profits from transshipment services down to zero. Markets for goods, land and labor clear in each city, and markets for transshipment services clear in each port city. Appendix B.1 provides a formal definition and characterization of the equilibrium.

4.2 City populations in the model

What determines the population of cities in equilibrium? The model delivers the following structural equation for the equilibrium population of city r, $N(r)$:

$$N(r)^{[1+\eta\sigma+(1-\gamma-\alpha)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} = \gamma^{\sigma-1}\tilde{a}(r)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} A(r)^{\frac{(\sigma-1)^2}{2\sigma-1}} (1-F(r))^{(1-\gamma)\frac{(\sigma-1)^2}{2\sigma-1}} MA(r) \quad (7)$$

where $MA(r)$ is the *market access* of city r, given by

$$MA(r) = \sum_{s=1}^{S} \frac{\tilde{a}(s)^{\frac{(\sigma-1)^2}{2\sigma-1}} A(s)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} (1-F(s))^{(1-\gamma)\frac{\sigma(\sigma-1)}{2\sigma-1}} N(s)^{[1-\eta(\sigma-1)-(1-\gamma-\alpha)\sigma]\frac{\sigma-1}{2\sigma-1}}}{\mathbf{E}[T(r,s)]^{\sigma-1}} \quad (8)$$

and $\tilde{a}(r)$ can be obtained by scaling amenities $a(r)$ according to

$$\tilde{a}(r) = \aleph_c a(r)$$

where the endogenous country-specific scaling factor $\aleph_c$ adjusts such that the exogenously given population of country c equals the sum of the populations of its cities.

Equation (7) implies that the population of city r is increasing in four objects: (1) re-scaled city amenities $\tilde{a}(r)$; (2) fundamental city productivity $A(r)$; (3) the share of land allocated to production, $1-F(r)$; (4) and the city's market access, $MA(r)$.

How is the population of a port city affected by the development of its port? The following proposition shows that the net effect on population is the outcome of three opposing forces: the *market access effect* that increases the population of the city, and the *land use* and *disamenity effects* that lead to a decrease in the city's population.

Proposition 1. *An increase in the share of land allocated to the port in city r, $F(r)$, decreases shipping costs $\mathbf{E}[T(r,s)]$, thus increasing $MA(r)$. Everything else fixed, an increase in $MA(r)$ increases the population of the city (market access effect). Holding $MA(r)$ fixed, an increase in $F(r)$ decreases the share of land that can be used for production, $1-F(r)$, thus decreasing the population of the city (land use effect). Finally, an increase in $F(r)$ increases shipping flows, thus lowering amenities $\tilde{a}(r)$ and, everything else fixed, city population (disamenity effect).*

Proof. These results follow directly from equation (7). □

Proposition 1 sheds light on the fact that, to measure the net effect of port development, it is essential to consider both its benefits and its costs. On the one hand, port development lowers shipping costs. On the other hand, it requires scarce local land that needs to be reallocated from other productive uses, while also making the city a less desirable place to live. The model, and equation (7) in particular, provide a structure that allows us to capture these opposing forces.

5 Empirical evidence for the model's mechanisms

Building on the model, we revisit the empirics in order to test the positive and negative effects of port development predicted by the model. We estimate the following panel specification, referred to as the *model-inspired empirical specification*:[28]

$$\ln(Pop_{it}) = \phi_1 * \ln(Ship)_{it} + \phi_2 * \ln(MA)_{it} + \alpha_i + \delta_t + \epsilon_{it} \tag{9}$$

where $\ln(MA)_{it} = \ln\left(\sum_{s=1}^{S} \frac{Pop_{st}^{[1-\eta(\sigma-1)-(1-\gamma-\alpha)\sigma]\frac{\sigma-1}{2\sigma-1}}}{T_t(i,s)^{\sigma-1}}\right)$ is the empirical equivalent of the model-based market access term, and all other variables are as previously defined in Section 3. According to the mechanism described in the model, we expect ϕ_1 to be negative and ϕ_2 to be positive.

To calculate trade costs $T_t(i,s)$, we use the fast marching algorithm to calculate the lowest overall shipping cost between any given pair of cities. To calculate the shipping costs across all potential routes, we need to calculate each possible component: i) the cost of shipping overland; ii) the cost of sea shipping; and iii) the cost of transshipment at seaports. Following Allen and Arkolakis (2014), we assume that overland shipping costs ϕ_ς and sea shipping costs ϕ_τ take the form

$$\phi_\varsigma(d) = e^{t_\varsigma d} \qquad\qquad \phi_\tau(d) = e^{t_\tau d}$$

where d is (point-to-point) distance traveled. We take the values of t_ς and t_τ from the road and sea shipping cost elasticities estimated by Allen and Arkolakis (2014).

To estimate transshipment costs across seaports (component iii) of trade costs), we take estimates of port costs from Blonigen and Wilson (2008). Consistent with increasing returns to scale in port activities (Rodrigue, 2016), we find a negative and statistically significant association between port costs and the size of shipping flows. We use the estimated coefficient from this regression to predict port costs for all the ports in our data for each decade.[29] Note that changing transshipment costs are the only source of time series variation in our estimated trade costs. The model-based measure of market access also requires taking a stand on the values of the parameters. Table D.9 contains the parameter values we use and their source.

Both regressors in the model-inspired specification (9) are potentially endogenous, requiring two sources of exogenous variation. We use depth as an instrument for shipping, as explained in

[28]Due to the lack of time-varying data on ports' land use, we cannot directly take equation (7) to the data.
[29]Appendix C.13 provides the full details and results of the estimation.

Section 3. In addition, we use an exogenous population-growth shifter based on regional climate to construct an instrument for market access. This IV is based on insights from the urban economics literature, which has found that people have moved to places with warm winters over the course of the 20th century (e.g., Oi (1996); Rappaport (2007)).

We use the average number of frost free days, $frostfree_i$, during the years between 1961-1990 in each city to predict population growth during our time period.[30] In order to predict population, we estimate the following specification:

$$\ln(Pop)_{it} = \sum_{k=1960}^{1990} \beta_k * frostfree_i * \mathbb{1}(Year = k) + \alpha_i + \delta_{ct} + \epsilon_{it}$$

where β_k estimates the effect of warmer winters on population in each decade, α_i denotes city-specific fixed effects, and δ_{ct} allows for the full set of country by year fixed effects. Inclusion of these implies that we only use *within-country* variation in climatic conditions when estimating the effect of frost-free days on population growth. We do this to address the concern that climatic conditions vary across regions in ways that may correlate with unobserved drivers of population growth, confounding our estimates of interest. Table D.10 shows the result of this estimation and presents some robustness checks. To construct our second instrument, we predict population for each city-year pair based on the estimated effects of frost free days and the estimated city fixed effect (we do not use the estimated country-year fixed effects to predict population). Using these predictions for city-level population, we define our second instrument as follows:

$$\ln\left(MAIV_{it}\right) = \ln\left(\sum_s \frac{\exp(\widehat{\ln(Pop)}_{it})}{\left(T_{1950}\left(i,s\right)\right)^{\sigma-1}}\right)$$

where $T_{1950}\left(i,s\right)$ is the trade cost between cities i and s in 1950. We hold bilateral trade costs fixed throughout all years in order to make sure that potentially endogenous changes in trade costs over time are not used in the instrument. The specifications are estimated on the set of *port* cities in our dataset. Importantly, however, the market access of port cities is calculated using the full set of (port and non-port) cities.

Table 3 presents the estimation results. The OLS estimate in column (1) shows a very small negative effect of shipping on population that is not distinguishable from zero. Column (2) shows the 2SLS specification. Consistent with the predictions of the model, once we control for market access, shipping has a negative, statistically significant effect on population.[31] The instruments

[30]Appendix C.14 contains a description of the data on the number of frost free days. This second instrument uses a source of exogenous variation that is orthogonal to port depth, as the number of frost free days and port depth are uncorrelated with each other. The correlation between the number of frost free days and depth is 0.04 (p-val: 0.40).

[31]As expected, market access has a significant positive effect on population. Jedwab and Storeygard

yield a combined Kleibergen-Paap F-statistic of 9.63 which is just below the often recommended value of 10; however, it is larger than the critical value of 7.03 that the Stock-Yogo weak ID test suggests for 10% maximum bias (Stock and Yogo, 2002). Columns (3) and (4) report the first stages of the regression. Reassuringly, depth is a strong predictor of shipping, while the market access IV predicts market access strongly. Table D.11 shows that the pre-trends check with respect to depth holds (for both first stages) in this more complex specification that adds market access.

We test the robustness of this result in a number of ways in Table D.12. First, we show that the results are remarkably robust to dropping cities in the close vicinity of the city in the market access IV, suggesting that much of the identifying variation is coming from population movements further away from the city itself. Moreover, the signs of the effects are robust to the same set of controls used in Section 3, though in the case of these demanding specifications, we don't always retain statistical significance at 10%. Finally, we show that the coefficients are robust to applying the Borusyak and Hull (2022) correction to potential non-random shock exposure in the market access IV by constructing counterfactual shocks that could have been realized.[32]

A potential interpretation of the negative effects that containerization induced on port city population is that they corresponded, mechanically, to job losses in the shipping sector. It is true that containerization is a labor-saving technology that may have led to job destruction as a result of automation. In Appendix B.9, we examine the extent to which post-1960 job losses in the U.S. water transportation sector can account for the population losses in port cities we obtain in Table 3. We find that job destruction in water transportation is unable to quantitatively account for more than a small fraction of our estimated population losses.

We conclude that the results in this section lend well-identified evidence to the model mechanism. In the next section, we therefore turn to taking the full model to the data.

6 The aggregate effects of containerization

We use our model to measure the aggregate effects of containerization. To this end, we first take the model to post-containerization data, then roll back containerization in a counterfactual exercise. The difference between the counterfactual and the post-containerization equilibrium captures the aggregate effects of containerization. We also test whether the model can replicate the reduced-form findings of Section 3.

(2022) are the only paper we are aware of that report standardized coefficients that allow for a comparison of magnitudes. They estimate that a one standard deviation increase in market access leads to a 0.43 – 0.85 standard deviation increase in population. Relative to that paper, our estimate is slightly larger (1.13), but within the same ballpark.

[32] To construct these counterfactual shocks, we reshuffle frost-free days across the world (column 7) or within 30 degree latitude bands (column 8).

6.1 Taking the model to the data

Taking the model to the data consists of three steps. First, we calculate inland and sea shipping costs across cities and choose a functional form for endogenous transshipment costs. Second, we choose the values of the model's eight structural parameters. Finally, we back out the values of unobserved city fundamentals that rationalize the post-containerization data.

6.1.1 Calculating shipping costs

To calculate inland and sea shipping costs across cities as a function of distance d, we follow the approach described in Section 5 based on Allen and Arkolakis (2014). Next, we choose endogenous transshipment costs as a function of the share of land allocated to transshipment services (*port share*, F), $\psi(F)$. Our goal is to keep the functional form of ψ numerically tractable and to satisfy our theoretical restrictions. One simple function that satisfies both is

$$\psi'(F) = 1 - F^{-\beta} \tag{10}$$

where we restrict $\beta > 0$ to guarantee $\psi' < 0$.

Finally, we capture the additional costs of cross-country trade, such as tariffs, quotas and red-tape barriers, by multiplying the overall shipping cost between any two cities that are not in the same country by a constant $B > 1$. We choose the value of B such that the model replicates the ratio of international trade to world GDP in 1990. This procedure yields $B = 2.1$.

6.1.2 Choosing the values of structural parameters

On the production side, we take the estimate of the strength of agglomeration externalities, $\alpha = 0.06$, from Ciccone and Hall (1993). The expenditure shares on labor and land equal γ and $1-\gamma$, respectively. We base our benchmark value of γ on Desmet and Rappaport (2017), who estimate a value of 0.10 for the difference between the land share and the agglomeration elasticity in the United States between 1960 and 2000, a period that corresponds to our sample period. Given we set $\alpha = 0.06$, this suggests choosing $\gamma = 0.84$.[33]

On the consumption side, we have three structural parameters: the migration elasticity, which we set to $\eta = 0.15$ based on Kennan and Walker (2011); the elasticity of substitution across tradable final goods, which we set to $\sigma = 4$ based on Bernard, Eaton, Jensen, and Kortum (2003); and the elasticity of port city disamenities with respect to shipping, which we set to $\rho = 0.005$ based on the estimated economic cost of pollution for Los Angeles from Marquez and Vallianatos (2012).[34]

Finally, there are three structural parameters that influence shipping costs. One is the dispersion

[33] Another advantage of using this land share estimate is that it also accounts for the share of land embedded in housing, which is absent from our model but could matter for the quantitative results.

[34] See section B.10 for details.

of idiosyncratic shipping costs, which – together with the functional form of these costs – we take from Allen and Arkolakis (2019), setting $\theta = 203$. Another is the elasticity of transshipment costs to total shipping at the port (congestion externalities), which we take from the empirical estimates of Abe and Wilson (2009), setting $\lambda = 0.074$. Table D.9 summarizes the calibration of these seven structural parameters.

The last structural parameter to choose is β from the endogenous transshipment function. Given that β drives the relationship between the value of shipping flows and the port share (see equation B.5 in Appendix B.2), we calibrate it to match the correlation between these two variables in the data.[35] Under higher values of β, the endogenous port development mechanism plays a stronger role in the model. Hence, everything else fixed, landlords have an incentive to increase the port share further if β is high. Thus, we expect a stronger correlation between shipping and port share under higher values of β. This is precisely what we find. Figure E.5a plots the values of the correlation for a range of β between 0.020 and 0.046. Within this range, $\beta = 0.031$ is the one that implies the correlation found in the data, 0.474.

6.1.3 Recovering post-containerization fundamentals

We use observed data on city populations, shipping flows and city-level GDP per capita together with the structure of the model to find the set of fundamental city amenities $\bar{a}(r)$, productivities $A(r)$ and exogenous transshipment costs $\nu(r)$ that rationalize the data.

As city-level GDP data are only available for 1990, we choose to back out the model fundamentals based on the 1990 distribution of population, shipping and GDP. Hence, the aggregate effect of containerization can be assessed by comparing the counterfactual equilibrium (pre-containerization) to our 1990 equilibrium (post-containerization).

We transform the number of ships observed in the data in port city r in 1990, $Ship(r)$, into the value of shipments, $Shipping(r)$, according to

$$Shipping(r) = V \cdot Ship(r)$$

where we choose V to match the ratio of shipping to world GDP. The rationale behind choosing this particular moment is that it can be calculated as a simple linear function of V:

$$\frac{\sum_r Shipping(r)}{\sum_r GDP(r)} = V \cdot \frac{\sum_r Ship(r)}{\sum_r GDP(r)}$$

where $Ship(r)$ and $GDP(r)$ are both observable in the data. This procedure gives us a value of

[35] The correlation between shipping and port share for the seven cities for which we have data is 0.474. Data sources are documented in Appendix C.5.

$V = 364$.[36]

Using city-level GDP data, we can obtain wages as

$$w\left(r\right) = \gamma \frac{GDP\left(r\right)}{N\left(r\right)}$$

according to the model, where the structural parameter γ is calibrated to 0.84.

Once population $N\left(r\right)$ and wages $w\left(r\right)$ are available for each city and the value of shipments, $Shipping\left(r\right)$, is available for each port city, the equilibrium conditions of the model can be inverted to back out city amenities up to a country-level scale, $\tilde{a}\left(r\right)$, fundamental city productivities $A\left(r\right)$, and each port city's exogenous transshipment costs $\nu\left(r\right)$. We provide the details of this inversion procedure in Appendix B.3.[37]

6.2 Counterfactual: rolling back containerization

In our counterfactual, we account for the technological aspects of containerization in seaports that we document in Sections 1 and 3: lower costs, particularly in deep ports, and the increased land-intensity of transshipment. This requires us to change the values of three model fundamentals relative to the post-containerization equilibrium. We discuss these three changes in model fundamentals next.[38]

First, we capture the *lower land intensity* of port technology before containerization by decreasing the shape parameter of transshipment technology, β. As we argued in Section 6.1.2, a decrease in β makes the endogenous transshipment cost function less responsive to changes in the port share, $F\left(r\right)$. Hence, under lower values of β, port city landlords have less incentive to increase $F\left(r\right)$, and port sizes will be generally smaller.

To choose the value of the parameter in the counterfactual, β_{CF}, we use the empirical evidence of Section 3.2. In particular, we found that containerized ports occupy on average 75% larger area if we hold the volume of traffic fixed. In our model, this means that the average port share would have increased by 75% if we held the non-technological determinants of the port share, i.e., shipping and land rents, fixed. Figure E.5b shows that the increase in mean port share is monotonic in β_{CF}, the counterfactual value of β, and hence the parameter is identified. The value at which mean port share increases by 75% is $\beta_{CF} = 0.018$.

[36] As not all our port cities have positive shipping flows in 1990 but the model cannot rationalize zero shipping flows under finite positive values of city-specific fundamentals, we change $Ship\left(r\right)$ from zero to one in these cities.

[37] The complex structure of the model does not allow us to prove that the inversion procedure identifies a unique set of $\tilde{a}\left(r\right), A\left(r\right)$ and $\nu\left(r\right)$. Nonetheless, we have experimented with various different initial guesses, and the inversion algorithm converges to the same fixed point, suggesting that the vector of city-specific fundamentals that rationalize the data is likely unique.

[38] Appendix B.4 provides details on how we numerically solve for counterfactual equilibria in the model.

Second, we capture the fact that *depth* was not relevant for transshipment prior to containerization. To this end, we offset the relationship between exogenous transshipment costs and depth in the counterfactual. We first run the regression

$$\log \nu(r) = \omega_0 - \omega_1 * Depth(r) + \varepsilon(r)$$

on our sample of port cities, where $\nu(r)$ is the exogenous transshipment cost of city r recovered in Section 6.1.3, and $Depth(r)$ is the depth measure, defined in Section 3. In line with the fact that depth lowers transshipment costs after containerization, we find $\widehat{\omega_1} = 0.048$ (s.e. 0.025, p-value 0.053). Next, we undo this dependence of exogenous transshipment costs on depth by adding $\widehat{\omega_1} * Depth(r)$ to $\log \nu(r)$.

Finally, we incorporate the overall *reduction in transshipment costs* due to containerization by increasing exogenous transshipment costs $\nu(r)$ uniformly across ports. More precisely, we increase $\log \nu(r)$ by the same number ν_{CF} at each port to match the estimated 25% average change in the sum of exogenous *and* endogenous transshipment costs as a result of containerization.[39] Naturally, higher values of ν_{CF} yield a larger change in transshipment costs, suggesting that there should be a unique ν_{CF} at which we meet our 25% target. This procedure identifies $\nu_{CF} = 0.245$.

Note that, in our counterfactual, we focus on estimating the effects containerization had by reducing transshipment costs at seaports. In reality, containerization had broader effects on transport costs. Most importantly, it arguably also reduced overland transport costs as intermodal transshipping between trucks and railways became cheaper. To address this, Appendix B.8 investigates how adding an additional inland cost increase to our counterfactual simulation changes the results.

6.3 Test of the model: the local effects of containerization

In this section, we test our quantified model by studying whether it can replicate the three stylized facts about the local effects of containerization that we documented in the empirics (Section 3). We summarize the results from Section 3, reporting standardized coefficients, in column (1) of Table 4. In column (2), we present the results from running similar regressions with data simulated from our baseline model. For example, to replicate Stylized fact 1 (panel A), we report estimates of the two-period version of regression equation 1, where the periods represent either the 1990

[39] Rodrigue (2016, p. 117) estimates that containerization led to an overall 70% to 85% reduction in maritime transport costs by 2010; "While before containerization maritime transport costs could account for between 5 and 10 percent of the retail price, this share has been reduced to about 1.5 percent, depending on the goods being transported." A reduction from 5% to 1.5% of retail price equals a 70% cost reduction ($= 1 - 1.5/5$); similarly, a reduction from 10% to 1.5% equals an 85% cost reduction. We estimate that 36% of the total cost reduction took place up to 1990, by assuming that cost reductions are proportionate to ship size increases. These calculations are based on data from the *Miramar Ship Index* (Haworth, 2020). More details on these data are provided in Appendix C.11. Using the more conservative estimate of 70%, this gives us a 25% decrease in average transshipment costs.

equilibrium ($After_t = 1$) or the no-containerization counterfactual ($After_t = 0$). As can be seen in column (2), the estimated effect of depth on shipping growth is positive and significant in our baseline model, just like in the data (column 1). This shows that the model can replicate Stylized fact 1, the disproportionate growth of shipping in deep port cities.

To examine whether Stylized fact 2 holds in the model, we consider the two-period version of our land scarcity result (specification 2). As our data-based land scarcity measure does not have a role in the model, we cannot use it in this specification. A model variable that naturally captures the local scarcity of land is land rents. Thus, we interact depth with log pre-containerization land rents, $\ln R_{i,CF}$, to see whether containerization reallocates shipping toward port cities with less land scarcity in the model. As column (2) in panel B shows, the effect of land rents on shipping growth is negative and statistically significant, just like the effect of our land scarcity measure in the data (column 1).

Finally, we examine the local population effects of shipping (Stylized fact 3) by estimating equation (3) in the model. Following the same identification strategy as in the empirics, we instrument the change in shipping with $Depth_i$. Column (2) in panel C presents the results of this exercise. The estimated coefficient is neither statistically, nor economically significant, as in the data (column 1). This result suggests that the local costs of port development, in the form of land use and disamenities, are sufficient to offset the positive market access effect not only in the data, but also in the model. Note that we did not target this zero effect in the model calibration.

To examine whether local port development costs drive these results, we contrast column (2) with a model in which port development does not entail local costs. In Appendix B.5, we develop a benchmark model in which, just like in our baseline model, port development yields transshipment cost reductions. We label this model 'Benchmark 1.'[40] In Benchmark 1, transshipment cost reductions are exogenous and *free*: they do not require land, and they do not induce local disamenities in the port city. We take Benchmark 1 to the data and conduct the same no-containerization counterfactual in it.

Column (3) of Table 4 shows how Benchmark 1 fares in replicating the three stylized facts of containerization. Three results emerge. First, Benchmark 1 also features disproportional shipping growth in deep ports (panel A). This is because we offset the relationship between transshipment costs and depth in Benchmark 1 just like in our baseline model. Second, the heterogeneous effect of containerization with respect to land scarcity turns insignificant (panel B). This is because port development no longer puts a strain on local land. Thus, in this type of model, port development takes place to the same extent in land-scarce as land-abundant cities. Finally, shipping leads to a significant increase in city population in Benchmark 1 (panel C). This is also intuitive: while better market access draws people into the port city, the offsetting forces (the local costs of port

[40]We develop another benchmark model, named 'Benchmark 2,' in Section 6.4.

development) are absent. This finding underscores that the local costs of port development drive the zero population effect of shipping in our baseline model.

To see how much land use and port city disamenities contribute to the difference between our baseline model and Benchmark 1, we shut down these two forces separately in columns (4) and (5) of Table 4. In column (4), we shut down land use but keep disamenities ("Benchmark 1 with disamenities"). This turns out to change Stylized Fact 2 and 3 in ways that contradict the data. The heterogeneous effect of land scarcity on shipping (Stylized Fact 2) is small and statistically indistinguishable from zero. While the inclusion of disamenities from port activities does lower the effect of shipping on population relative to Benchmark 1 as we would expect, the coefficient remains statistically significant (at 10%) and larger than in the data.

In column (5), we shut down disamenities but assume that port development uses local land, as in our baseline model ("Baseline model without disamenities"). In this case, the results are closer to column (2): land scarcity leads to significantly lower shipping growth, and the population effects of shipping turn insignificant, though the point estimate of the latter effect is larger than in column (2). Overall, these results suggest that the land use effect is particularly important to explain our empirical facts, but port city disamenities also play a role in offsetting the market access effect on local population.

6.4 The aggregate effects of containerization

We estimate that aggregate world welfare increased by 3.23% as a result of containerization.[41] Rolling back containerization increased the international trade to world GDP ratio by 4.1 percentage points from the counterfactual to the 1990 equilibrium. As a reference point, the trade to world GDP ratio increased by 15 percentage points between 1960 and 1990. This suggests that containerization was responsible for about *one quarter* of the overall increase in trade to world GDP during these three decades.

The fraction of land occupied by ports (i.e., the port share) increases in most port cities from the counterfactual to the 1990 equilibrium. Port shares become larger for two reasons. First, the increase in β increases the incentive to invest more land in port development. Second, the reduction in trade costs leads to increased demand for shipping, encouraging yet more investment in port development. Figure E.6 presents the full distribution of port share changes across cities. The median change is 5 percentage points, while the 5th percentile is zero pp and the 95th percentile is 46 pp.

How large was the cost of increased land use due to containerization? To answer this question,

[41] We define the change in aggregate world welfare as the average of changes in country-level welfare between the counterfactual and the 1990 equilibrium, weighted by country population. Within each country, labor mobility equalizes welfare across cities, as in Redding (2016). However, we do not allow for mobility across countries, hence different countries experience different welfare effects.

we conduct a decomposition that exploits the fact that the welfare gains from containerization stem from a combination of three factors in the model. First, containerization lowers shipping costs, thus increasing welfare. Second, containerization increases port city land use, which we can label as the *resource costs* of containerization. Finally, containerization might yield gains from increased specialization of cities in port or non-port activities, which we can label as the *specialization gains* from containerization.

To assess the quantitative importance of each of these margins, we leverage the benchmark model, 'Benchmark 1,' developed in Section 6.3. Here, the welfare gains from containerization only stem from shipping cost reductions. Next, we develop an additional benchmark model, which we label as 'Benchmark 2.' Benchmark 2 requires land to be used to reduce transshipment costs. However, we restrict land use to be identical across port cities (and equal to the mean port share in our baseline).[42]

As Benchmark 2 only differs from Benchmark 1 in land being used for port activities, a comparison between these two models reveals the resource costs of increased land use due to containerization. As our baseline model only differs from Benchmark 2 in the potential specialization of port cities in port or non-port activities (through each city choosing the allocation of land between the two), a comparison between these two models reveals the endogenous specialization gains from containerization.

We find that containerization leads to welfare gains of 3.48% in Benchmark 1. In Benchmark 2, the gains from containerization reduce to 2.84%. The difference between Benchmark 1 and Benchmark 2, 0.64 percentage points, captures the resource costs of containerization. These costs are sizeable: they account for as much as 18% of the gains from the shipping cost reduction. Finally, the difference between Benchmark 2 and our baseline model, 0.40 percentage points, captures the specialization gains from containerization. Note that these gains are able to offset about 63% of the resource costs of containerization, but they do not fully compensate for all the costs.

In Appendix B.8, we show that the aggregate and local effects of containerization implied by the model are robust to different values of the containerization shock and some alternative modeling choices. These alternative specifications include different values of transshipment cost shape parameter β in 1990 or in the counterfactual, different changes in exogenous transshipment costs, and a model in which landlords make profits from the provision of transshipment services.

[42] As our goal is to estimate the resource cost of land use separately from the cost of disamenities, we use the version of Benchmark 1 with disamenities in the decomposition exercise. Similarly, Benchmark 2 features disamenities as well. We provide a detailed description of each benchmark model and their quantitative estimation in Appendix B.5.

7 The effects of targeted port development

We use our estimated model to illustrate the effects of targeted port development policy in this section. We study a large-scale port development policy similar to the Chinese government's Maritime Silk Road project, which is part of the 'Belt and Road Initiative.'[43] In particular, we study the effects of a 10% reduction in exogenous transshipment costs in 24 port cities in Asia, Africa and Europe targeted by Chinese investment (see Figure E.7 for the set of targeted ports).[44]

Table D.15 examines the effects of this policy on treated and untreated port cities, and inland cities. We compare the effects generated by our model ('Baseline') to those of a more standard model ('Benchmark 1' – introduced in Section 6.3). As column (1) demonstrates, targeted port cities see a significant and large increase in shipping activities, primarily at the expense of non-targeted port cities in the same country. This local reallocation of shipping is more pronounced in the baseline model than in Benchmark 1 (column 5). To see why this is the case, in columns (2) and (6) we examine the effect on port costs (the sum of exogenous and endogenous transshipment costs, $\nu(r) + \psi(F(r))$). In Benchmark 1, endogenous transshipment costs are absent, implying that targeted port cities see an exact 10% (0.105 log point) decline in their transshipment costs, while non-targeted cities see no effect. By contrast, in the baseline model, the direct effect of the policy is amplified by an endogenous reallocation of land within the city. This results in a decline in endogenous transshipment costs in targeted ports (where more land is allocated to the port) and an increase in endogenous transshipment costs in non-targeted ports (where less land is allocated to the port). This endogenous port development response to the policy is precisely what draws additional shipping into targeted cities and away from non-targeted ones.[45]

We also study the effects on cities' market access (as defined by equation 8) and population across both models. The effect on market access is similar in both simulations. In terms of population responses, however, the similar improvement in market access results in strikingly different population responses – highlighting the local costs of port development at work in our model. In the baseline (column 4), endogenous port development in targeted port cities moves people out of the city through increased land use and disamenities, primarily to non-targeted port cities. In contrast, in Benchmark 1 (column 8), targeted ports gain population.

We examine how targeted port development redistributes shipping and real GDP across regions of the world in Figure 2. We find the most dramatic distributional effects in Asia. Strikingly, we see

[43]The simulation we conduct is *similar* to the Maritime Silk Road project, as we analyze effects relative to the 1990 equilibrium, not today. Moreover, the absence of specific details on the size of the actual investments precludes us from matching exactly what the project entails.

[44]We take the targeted ports from OECD (2018) and choose the decrease in $\nu(r)$ to be 10% to illustrate the effects of a sizeable, but not dramatic decrease in transshipment costs. We keep all other fundamentals of the model fixed at their levels recovered in Section 6.1.

[45]As presented in Table D.16, the results remain similar if we include country fixed effects.

a dramatic reallocation of shipping to China and away from Singapore (which we estimate loses almost 50% of its shipping flows).[46] Neither countries have targeted ports in this simulation. While these effects are also present in the benchmark model, they are far more muted. In our model, the initial reallocation of shipping is amplified by endogenous port development. Interestingly, Singapore sees a more than 1% gain in real GDP in our baseline model, as the city's declining port frees up land that can be used profitably outside the shipping sector. This is particularly true in the case of Singapore, where the non-port sector is very productive.[47] As the example of Singapore illustrates, endogenous port development has the potential to substantially amplify changes in shipping and real GDP in our baseline model relative to a standard trade model such as Benchmark 1.

8 Conclusion

The containerization shock studied in this paper allows us to shed light on the economic effects of port development. Our findings suggest that the land-intensive nature of port development is an empirically strong force that matters for the local, aggregate and distributional effects of port development. Recent disruptions to supply chains due to the COVID-19 pandemic have highlighted some of the consequences of these forces. The containers flowing out of the port of Long Beach in 2022 due to a lack of storage space suggest that many ports operate with very little slack capacity, limiting their ability to adjust to shocks.[48] Our analysis suggests that the scarcity of land around many of the world's major ports is an important driving force. In light of this, our findings raise the question of whether ports today are located optimally. Should highly productive cities such as Los Angeles or Singapore continue to specialize heavily in port activities? We leave the exploration of this normative question for future research.

References

Abe, K. and J. Wilson (2009). *Weathering the Storm: Investing in Port Infrastructure to Lower Trade Costs in East Asia*. World Bank.

Allcott, H. and D. Keniston (2017). Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America. *Review of Economic Studies 85*(2), 695–731.

Allen, T. and C. Arkolakis (2014). Trade and the Topography of the Spatial Economy. *Quarterly*

[46] It must be noted, however, that China's percentage change in shipping does not correspond to a dramatic absolute change, as China had relatively little shipping back in 1990.

[47] According to our model, Singapore is at the 98th percentile in the world productivity distribution. Of course, the economic benefits from dismantling a port may be not the only factor considered by decision-makers in reality. Governments' objective functions may include geopolitical advantages from maintaining a central position in the global shipping network. In our analysis, we focus on the economic effects and do not consider these additional factors.

[48] E.g., https://qz.com/2079345/cargo-ships-containers-are-piling-up-in-long-beach.

Journal of Economics 129(3), 1085–1140.
Allen, T. and C. Arkolakis (2019). The Welfare Effects of Transportation Infrastructure Improvements.
Allen, T. and D. Atkin (2022). Volatility and the Gains from Trade. *Econometrica 90*(5), 2053–2092.
Altomonte, C., I. Colantone, and L. Bonacorsi (2018). Trade and Growth in the Age of Global Value Chains. Technical report, BAFFI CAREFIN Centre Research Paper.
Anderson, J. (1979). A Theoretical Foundation for the Gravity Equation. *American Economic Review 69*(1), 106–116.
Armenter, R., M. Koren, and D. Nagy (2014). Bridges.
Bernard, A., J. Eaton, J. Jensen, and S. Kortum (2003). Plants and Productivity in International Trade. *American Economic Review 93*(4), 1268–1290.
Bernhofen, D., Z. El-Sahli, and R. Kneller (2016). Estimating the Effects of the Container Revolution on World Trade. *Journal of International Economics 98*, 36–50.
Bleakley, H. and J. Lin (2012). Portage and Path Dependence. *Quarterly Journal of Economics 127*, 587–644.
Blonigen, B. and W. Wilson (2008). Port Efficiency and Trade Flows. *Review of International Economics 16*(1), 21–36.
Borusyak, K. and P. Hull (2022). Non-Random Exposure to Exogenous Shocks. *National Bureau of Economic Research* (Working Paper 27845).
Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2020). Geography, transportation, and endogenous trade costs. *Econometrica 88*(2), 657–691.
Bridgman, B. (2021). Why Containerization Didn't Reduce Ocean Shipping Costs, At First. *mimeo*.
Brinkman, J. and J. Lin (2022). Freeway Revolts! The Quality of Life Effects of Highways. *Review of Economics and Statistics*, forthcoming.
Brooks, L., N. Gendron-Carrier, and G. Rua (2021). The Local Impact of Containerization. *Journal of Urban Economics 126*(C).
Ciccone, A. and R. Hall (1993). Productivity and the Density of Economic Activity. *National Bureau of Economic Research* (Working Paper 4313).
Coşar, A. and P. Fajgelbaum (2016). Internal Geography, International Trade, and Regional Specialization. *American Economic Journal: Microeconomics 1*(8), 24–56.
Corbett, M. (2010). *The History and Transformation of the Port of San Francisco, 1848-2010*. San Francisco Architectural Heritage.
Corden, W. and J. Neary (1982). Booming sector and de-industrialisation in a small open economy. *Economic Journal 92*(368), 825–848.
Coşar, A. and B. Demir (2018). Shipping Inside the Box: Containerization and Trade. *Journal of International Economics 114*, 331–345.
Desmet, K. and J. Rappaport (2017). The Settlement of the United States, 1800–2000: The Long Transition Towards Gibrat's Law. *Journal of Urban Economics 98*, 50–68.
Donaldson, D. and R. Hornbeck (2016). Railroads and American Economic Growth: A "market access" Approach. *Quarterly Journal of Economics 131*(2), 799–858.
Donaldson, D. and A. Storeygard (2016). The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives 30*(4), 171–98.
Ducruet, C., S. Cuyala, and A. E. Hosni (2018). Maritime Networks as Systems of Cities: The

Long-term Interdependencies Between Global Shipping Flows and Urban Development. *Journal of Transport Geography 66*, 340–355.

Ducruet, C. and T. Notteboom (2023). *Port Systems in Global Competition Spatial-Economic Perspectives on the Co-Development of Seaports*. Routledge.

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

Eyre, J. (1964). Shipping Containers in the Americas. In *Pan American Union: Recent Developments in the Use and Handling of Unitized Cargoes*, pp. 38–42.

Fajgelbaum, P. and S. Redding (2022). Trade, Structural Transformation and Development: Evidence from Argentina, 1869-1914. *Journal of Political Economy 130*(5), 1249–1318.

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

Falvey, R. E. (1976). Transport Costs in the Pure Theory of International Trade. *The Economic Journal 86*(343), 536–550.

Fujita, M. and T. Mori (1996). The Role of Ports in the Making of Major Cities: Self-agglomeration and Hub-Effect. *Journal of Development Economics 49*, 93–120.

Ganapati, S., W. F. Wong, and O. Ziv (2022). Entrepôt: Hubs, Scale, and Trade Costs.

Gilman, S. (1983). *The Competitive Dynamics of Container Shipping*. Gower Publishing Company.

Gomtsyan, D. (2016). Rise of the Machines: Evidence from the Container Revolution.

Haworth, R. B. (2020). Miramar ship index.

Head, K. and T. Mayer (2014). *Gravity Equations: Workhorse, Toolkit, and Cookbook*, Volume 4 of *Handbook of international economics*. Elsevier.

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

Holmes, T. and E. Singer (2018). Indivisibilities in Distribution. *National Bureau of Economic Research* (Working Paper 24525).

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

Jedwab, R. and A. Storeygard (2022). The Average and Heterogeneous Effects of Transportation Investments: Evidence from Sub-Saharan Africa 1960-2010. *Journal of the European Economic Association 20*(1), 1–38.

Kahveci, E. (1999). Fast Turnaround Ships and Their Impact on Crews.

Kennan, J. and J. Walker (2011). The Effect of Expected Income on Individual Migration Decisions. *Econometrica 79*(1), 211–251.

Koenig, P., Y. Pigné, S. Poncet, M. Sanch-Maritan, and C. Duvallet (2023). Sold to china: Container traffic in the port of piraeus. Technical report.

Krugman, P. (1987). The Narrow Moving Band, the Dutch Disease, and the Competitive Consequences of Mrs. Thatcher: Notes on Trade in the Presence of Dynamic Scale Economies. *Journal of Development Economics 27*(1-2), 41–55.

Krugman, P. (2011). Comparative Advantage, Growth, And The Gains From Trade And Globalization: A Festschrift in Honor of Alan V Deardorff. Technical report, Citigroup Foundation Special Lecture.

Levinson, M. (2010). *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger*. Princeton University Press.

Marquez, J. N. and M. Vallianatos (2012, January). Importing Harm: U.S. Ports' Impacts on Health and Communities. Technical report, Trade, Health and Environment Impact Project.

Martín-Antón, M., d. C. J. M. Negro, V., J. S. López-Gutiérrez, and M. D. Esteban (2016). Review of coastal land reclamation situation in the world. *Journal of Coastal Research 75*, 667 – 671.

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica 71*(6), 1695–1725.

Monte, F., S. Redding, and E. Rossi-Hansberg (2018). Commuting, Migration, and Local Employment Elasticities. *American Economic Review 108*(12), 3855–90.

Moriconi-Ebrard, F. (1994). *Geopolis*.

Nagy, D. (2022). Hinterlands, City Formation and Growth: Evidence from the U.S. Westward Expansion. *Review of Economic Studies*, forthcoming.

OECD (2014). The Competitiveness of Global Port Cities. Technical report, OECD Publishing.

OECD (2018). The Belt and Road Initiative in the Global Trade, Investment, and Finance Landscape. Technical report, OECD Publishing Paris.

Oi, W. (1996). The Economics of New Goods. In *The Welfare Implications of Invention*. University of Chicago Press.

Pines, D. and E. Sadka (1985). Zoning, First-best, Second-best, and Third-best Criteria for Allocating Land for Roads. *Journal of Urban Economics 17*(2), 161–173.

Port of San Francisco (1971). San Francisco Port Needs, Shipping and Area Requirements. Technical report, Port of San Francisco.

Rappaport, J. (2007). Moving to Nice Weather. *Regional Science and Urban Economics 37*(3), 375–398.

Redding, S. (2016). Goods Trade, Factor Mobility and Welfare. *Journal of International Economics 101*, 148–167.

Redding, S. and M. Turner (2015). *Transportation Costs and the Spatial Organization of Economic Activity*, Volume 5 of *Handbook of regional and urban economics*, pp. 1339–1398. Elsevier.

Rodrigue, J. (2016). *The Geography of Transport Systems*. Taylor & Francis.

Rua, G. (2014). Diffusion of Containerization.

Saiz, A. (2010). The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics 125*(3), 1253–1296.

Santamaría, M. (2022). The Gains from Reshaping Infrastructure: Evidence from the division of Germany.

Solow, R. (1972). Congestion, Density and the Use of Land in Transportation. *The Swedish Journal of Economics 74*(1), 161–173.

Solow, R. and W. Vickrey (1971). Land Use in a Long Narrow City. *Journal of Economic Theory 3*(4), 430–447.

Stock, J. and M. Yogo (2002). Testing for Weak Instruments in Linear IV Regression. In *Essays in Honor of Thomas Rothenberg*. Cambridge University Press.

The Economist (2023). Investments in ports foretell the future of global commerce. *The Economist*.

UNCTAD (2019). Review of Maritime Transport. Technical report, United Nations Publication.

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

A Tables

Table 1: Depth predicts shipping flows, but only after 1960 (Stylized fact 1)

	Dependent variable: ln(Ship)				
Independent variables	(1)	(2)	(3)	(4)	(5)
Depth × post 1970					0.247***
					*0.131****
					(0.059)
					{0.052}
Depth × 1960	-0.051	0.029	0.050	-0.055	
	(0.063)	(0.069)	(0.066)	(0.068)	
Depth × 1970	0.222***	0.233***	0.278***	0.213***	
	(0.069)	(0.077)	(0.082)	(0.071)	
Depth × 1980	0.188**	0.212**	0.291***	0.192**	
	(0.079)	(0.085)	(0.090)	(0.081)	
Depth × 1990	0.255***	0.222**	0.312***	0.283***	
	(0.086)	(0.087)	(0.099)	(0.087)	
Observations	2765	2765	2765	2360	2765
R-squared	0.126	0.248	0.131	0.142	0.126
Number of cities	553	553	553	472	553
Year FE	✓	✓	✓	✓	✓
City FE	✓	✓	✓	✓	✓
Population 1950 × Year	✓	✓	✓	✓	✓
Coastline × Year FE	✗	✓	✗	✗	✗
Land scarcity × Year	✗	✗	✓	✗	✗
GDP pc (country) × Year	✗	✗	✗	✓	✗

Notes: 'Depth' indicates the port suitability measure. It is interacted with decade dummies or an indicator variable for decades including and after 1970, as indicated. Standardized coefficient in italics underneath the baseline coefficient. Standard errors clustered at the city level in parentheses, Conley standard errors to adjust for spatial correlation in curly brackets. *** p<0.01, ** p<0.05, * p<0.1 (significance refers to clustered standard errors).

Table 2: The local causal effect of shipping on population (Stylized fact 3)

	Δln(Pop)	Δln(Pop)	Δ ln(Ship)	Δln(Pop)
Independent variables	(1)	(2)	(3)	(4)
Δln(Ship)	0.013	0.006		
	0.052	*0.022*		
	(0.009)	(0.073)		
	{0.014}	{0.115}		
Depth			0.272***	0.002
			*0.134****	*0.003*
			(0.086)	(0.020)
Observations	531	531	531	531
Specification	OLS	2SLS	FS	RF
KP F-stat		9.98		

Notes: 'Depth' indicates the port suitability measure. Standardized coefficients in italics underneath the baseline coefficients. All regressions control for population in 1950. Column (2) uses depth as IV for shipping. Notation for specification as follows: 'FS' refers to the first stage, 'RF' to the reduced form. Standard errors clustered at the city level in parentheses, Conley standard errors to adjust for spatial correlation in curly brackets. *** $p<0.01$, ** $p<0.05$, * $p<0.1$ (significance refers to clustered standard errors).

Table 3: Model-inspired specification: Disentangling the local benefits and costs of containerization

	(1)	(2)	(3)	(4)
Independent variables	ln(Population)	ln(Population)	ln(Ship)	ln(Market Access)
ln(Ship)	-0.001	-0.159**		
	(0.006)	(0.065)		
	{0.005}	{0.051}		
ln(Market Access)	1.512***	7.103***		
	(0.536)	(0.795)		
	{0.317}	{0.854}		
Depth × post 1970			0.275***	0.007***
			(0.058)	(0.001)
			{0.051}	{0.001}
Market Access IV			7.188	1.927***
			(5.428)	(0.140)
			{5.748}	{0.188}
Observations	2696	2696	2696	2696
Number of cities	544	544	544	544
Year FE	✓	✓	✓	✓
City FE	✓	✓	✓	✓
Population 1950 × Year	✓	✓	✓	✓
Specification	OLS	2SLS	FS	FS
KP F-stat		9.63		

Notes: 'Depth' indicates the port suitability measure. It is interacted with an indicator variable for decades including and after 1970. 'ln(Market Access)' is the empirical counterpart of the market access term, defined in Section 5. 'Market access IV' is the instrument for the market access term, defined in Section 5. 'FS' refers to the first stage. Standard errors clustered at the city level in parentheses, Conley standard errors to adjust for spatial correlation in curly brackets. *** p<0.01, ** p<0.05, * p<0.1 (significance refers to clustered standard errors).

Table 4: Stylized facts in model and data

	(1)	(2)	(3)	(4)	(5)
				Benchmark 1	Baseline model
	Data	Baseline	Benchmark 1	with disam.	w/o disam.
Panel A: Effect of depth on shipping (Stylized fact 1)					
	ln(Ship)	ln(Ship)	ln(Ship)	ln(Ship)	ln(Ship)
Depth	0.131***	0.262***	0.250***	0.249***	0.262***
	(0.031)	(0.011)	(0.010)	(0.010)	(0.011)
Observations	2,765	1,062	1,062	1,062	1,062
Panel B: Heterogeneous effect of land scarcity on shipping (Stylized fact 2)					
	ln(Ship)	ln(Ship)	ln(Ship)	ln(Ship)	ln(Ship)
Land scarcity * depth	-0.211**	-0.013**	-0.007	-0.007	-0.013**
	(0.096)	(0.005)	(0.005)	(0.005)	(0.005)
Observations	2,765	1,062	1,062	1,062	1,062
Panel C: Effect of shipping growth on population growth (Stylized fact 3)					
	Δln(Pop)	Δln(Pop)	Δln(Pop)	Δln(Pop)	Δln(Pop)
Δln(Ship)	0.022	-0.030	0.191***	0.095*	0.059
	(0.284)	(0.048)	(0.049)	(0.049)	(0.049)
Observations	531	531	531	531	531
First stage F-stat	9.98	584.5	680.3	679.9	585.6

Notes: All coefficients are standardized. Column (1) in Stylized Fact 1 repeats column (5) in Table 1; column (1) in Stylized Fact 2 repeats column (2) in Table D.5; column (1) in Stylized Fact 3 repeats column (2) of Table 2. The remaining columns use model-simulated data in equivalent regressions. *** p<0.01, ** p<0.05, * p<0.1.

B Figures

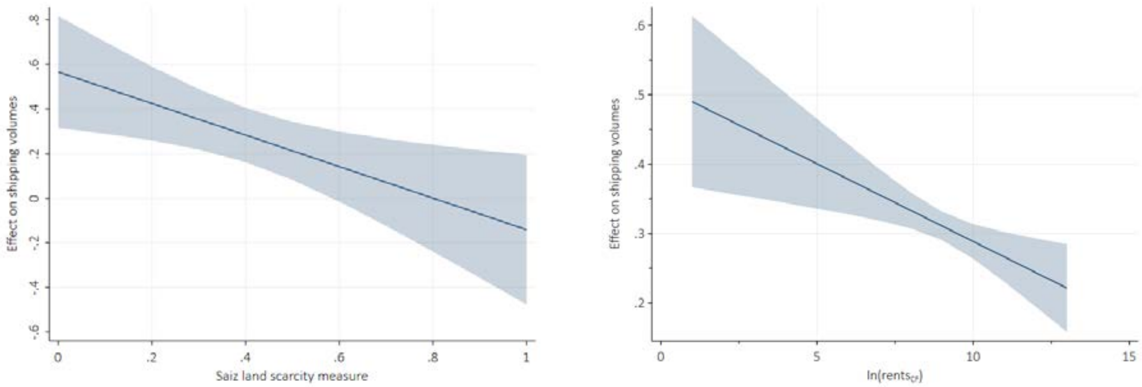

(a) Data

(b) Model-simulated data

Figure 1: Containerization increased shipping more where land is less scarce

Notes: Panel A shows the estimated γ coefficient from equation (2) evaluated at different values of the Saiz land scarcity measure. Panel B shows the same estimated coefficient using model-simulated data evaluated at different values of land rents in the no-containerization counterfactual.

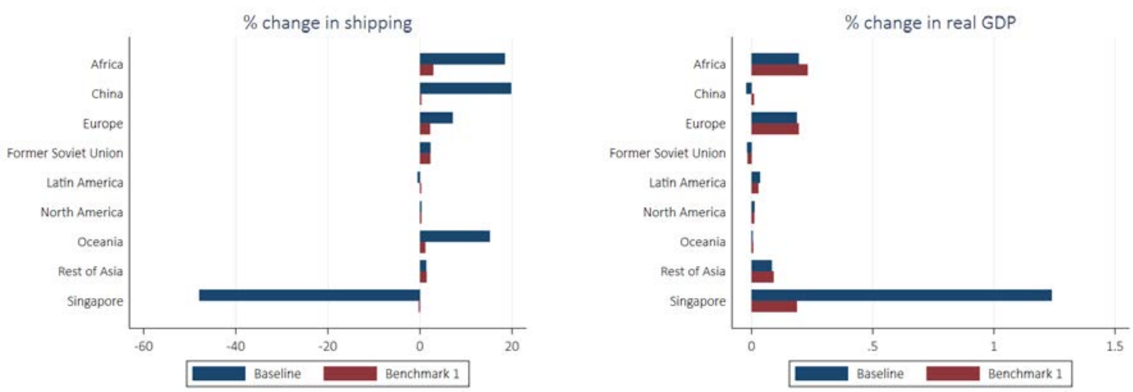

(a) Change in shipping, percent

(b) Change in real GDP, percent

Figure 2: Simulated changes across regions, Maritime Silk Road

Notes: Panel A (B) shows the change in total shipping (total real GDP) of each region between the model inversion and the Maritime Silk Road counterfactual. When delineating these regions, we roughly follow the world's continents. An exception is 'Rest of Asia,' which is Asia except China, Singapore, and the former Soviet Union. We treat China separately as we are naturally interested in the effects that this Chinese government policy has on China itself. We treat Singapore separately as we find strikingly large effects on this port city, which we discuss in the text.