Journal of International Economics 151 (2024) 103963

Contents lists available at ScienceDirect

Journal of International Economics

journal homepage: www.elsevier.com/locate/jie

All aboard: The effects of port development[☆]

César Ducruet [a], Réka Juhász [b,c,d,*], Dávid Krisztián Nagy [e,f,g,d], Claudia Steinwender [h,d]

[a] CNRS, 200 Avenue de la République, 92100 Nanterre, France
[b] UBC, 6000 Iona Drive Vancouver, BC V6T 1L4, Canada
[c] NBER, USA
[d] CEPR, France
[e] CREI, Carrer Ramon Trias Fargas 25-27, Barcelona, 08005, Spain
[f] UPF, Spain
[g] BSE, Spain
[h] LMU Munich, Ludwigstrasse 33, Munich, 80539, Germany

ARTICLE INFO

Dataset link: Replication for Ducruet, Juhasz, Nagy and Steinwender "All aboard: The effects of port development (Original data)

JEL classification:
R40
O33
F6

Keywords:
Ports
Containerization
Quantitative economic geography
Endogenous trade costs

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. Net of the costs, our results suggest that containerization in seaports increased world welfare by 3.4%.

[☆] We thank the editor, Costas Arkolakis, and two anonymous referees for comments that substantially improved the paper. We also thank Treb Allen, Leah Brooks, Don Davis, Dave Donaldson, Joseph Doyle, Nicolas Gendron-Carrier, Matt Grant, David Hummels, Giampaolo Lecce, Giacomo Ponzetto, Steve Redding, Roberto Rigobon, Esteban Rossi-Hansberg, Daniel Sturm, Tavneet Suri and Jaume Ventura for helpful comments and discussions. We thank Bruce Blonigen, Mario Martín Antón and their co-authors for kindly sharing data. Olalekan Bello, Sabrina Chen, Naman Garg, Yi Jie Gwee, Hamza Husain, Felix Iglhaut, Rodrigo Martínez Mazza, Emanuela Migliaccio, Verónica C. Perez, Shuhua Si, Yue Yu, Howard Zihao Zhang and a team of Columbia University undergraduate students provided outstanding research assistance. We gratefully acknowledge funding from the ERC and 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 (PID2019-111691RA-I00 and RYC2019-027620-I/AEI/10.13039/501100011033) and from the Generalitat de Catalunya, Spain, through CERCA and SGR Programme (2017-SGR-1393). César gratefully acknowledges financial funding from the French National Research Agency (ANR) project "MAGNETICS" No. ANR-22-CE22-0002 (Maritime Globalization, Network Externalities, and Transport Impacts on Cities).

[*] Corresponding author at: UBC, 6000 Iona Drive Vancouver, BC V6T 1L4, Canada.

E-mail addresses: cesar.ducruet@economix.fr (C. Ducruet), reka.juhasz@ubc.ca (R. Juhász), dnagy@crei.cat (D.K. Nagy), claudia.steinwender@econ.lmu.de (C. Steinwender).

https://doi.org/10.1016/j.jinteco.2024.103963
Received 13 July 2023; Received in revised form 7 June 2024; Accepted 11 June 2024

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

1. 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 related 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 (Ducruet et al., 2024). 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 et al., 2021; Altomonte et al., 2018). We construct a novel measure of 'naturally 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 narrative evidence 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.

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. Our estimates suggest that containerization increased world welfare by 3.4%, the ratio of world trade to GDP by 4.2 percentage points, and the median port size relative to city area by 2 percentage points. In a model-based decomposition, we find that the aggregate resource cost of increased land use amounted to 0.28% of world GDP, reducing the welfare gains from containerization by 8%. 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 45% 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

---

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

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 et al., 2014; Nagy, 2023). 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 et al. (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 et al., 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 container shipping networks in world trade (Wong, 2022; Heiland et al., 2023; Ganapati et al., 2022; Koenig et al., 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 3 discusses the main data sources used in the analysis. Section 4 presents three stylized facts about the local effects of containerization, while Section 5 introduces the model. Section 6 takes the model to the data, while Section 7 uses the quantified model to measure the aggregate effects of containerization and to illustrate the effects of targeted port development policies similar to the Maritime Silk Road. Finally, Section 8 concludes.

2. 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 F.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),

---

[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.

[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.

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 F.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.

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 h 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 h 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%.

3. 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 E.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 D).

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

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 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 D.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.

---

[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).

[6] These calculations are based on data from the *Miramar Ship Index* (Haworth, 2020). See Appendix D.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. (2023).

Saiz land scarcity measure (Appendix D.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 D.4). Data on city-level income levels are needed for the 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 D.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.

4. Stylized facts

In this section, we document three stylized facts about the local effects of containerization on port cities. Together, these stylized facts suggest that containerization entailed both costs and benefits for host cities.

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

Section 2 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 ft following Brooks et al. (2021). These authors argue that given vessel sizes in the 1950s (pre-containerization), depth beyond 30 ft conferred no advantage to the port. 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. 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 and validate this assumption in Appendix B.2 using nautical maps that show dredged channels.

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

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

The outcome variable of interest, $\ln(Ship_{it})$, is the log of shipping flows observed in city i at time t.[9] $Depth_i$ is the cross-sectional measure of port suitability defined above. We 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 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. We cluster standard errors at the city level in the baseline to account for the serial correlation of shocks.[10] 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.

---

[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.

[9] In practice, we replace the zeros in the data with ones and take the natural logarithm of this adjusted count (see Appendix B.1 for details).

[10] We also estimated Conley standard errors, but as these are typically very close to the clustered standard errors, we omit them for readability of the tables.

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

Independent variables	Dependent variable: ln(Ship)				
	(1)	(2)	(3)	(4)	(5)
Depth × post 1970					0.247*** _0.131_*** (0.059)
Depth × 1960	−0.051 (0.063)	0.029 (0.069)	0.050 (0.066)	−0.055 (0.068)	
Depth × 1970	0.222*** (0.069)	0.233*** (0.077)	0.278*** (0.082)	0.213*** (0.071)	
Depth × 1980	0.188** (0.079)	0.212** (0.085)	0.291*** (0.090)	0.192** (0.081)	
Depth × 1990	0.255*** (0.086)	0.222** (0.087)	0.312*** (0.099)	0.283*** (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. *** p<0.01, ** p<0.05, * p<0.1.

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. The results are also robust to allowing for regional trends (column 2 adds coastline-by-year fixed effects[11]), to allowing for differential trends across more and less land-scarce port cities (column 3 includes the Saiz land scarcity measure interacted with year fixed effects), and differential trends across initially rich and poor countries (column 4 adds country GDP per capita in 1960 interacted with year indicators).

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. Appendix B.3 discusses additional robustness checks. We note that the coefficient of interest 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. We now turn to examining whether this containerization-induced shipping boom was heterogeneous across port cities.

4.2. 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 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.[12]

---

[11] 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.'

[12] See Appendix A for further historical evidence and Appendix D.6 for a discussion of Seattle's containerization, respectively.

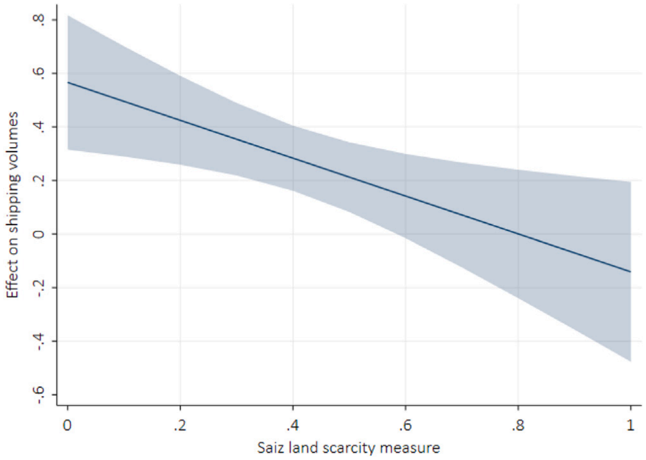

Fig. 1. Containerization increased shipping more where land is less scarce (Stylized fact 2).
Notes: This figure shows the estimated γ coefficient from Eq. (2) evaluated at different values of the Saiz land scarcity measure.

The land intensity of containerized terminals is also evident in contemporaneous data from *Google Earth*. Table E.2 shows that ports that handle more containerized cargo are typically larger. This is true when controlling for the total volume of traffic, and the results are also robust to the addition of other controls for cargo composition, and host country characteristics.

Port development took place where land was less scarce. The land-intensity of containerization documented above suggests the technology was better suited to locations where land for the expansion of the port was more readily available. To test this, we examine whether shipping increased more in cities where land was less scarce by allowing for heterogeneous effects with respect to land scarcity in regression Eq. (1):

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

where $LandScarcity_i$ measures the share of land in city i that cannot be built on, as defined in Section 3, 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 Fig. 1 (the corresponding estimates are presented in Table E.3). 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. Appendix B.3 discusses further robustness checks. This includes examining whether the Saiz land scarcity measure may be mismeasured due to land reclamation, which we find no evidence of.

A second test of whether the increased land requirements of port development affected the location of ports comes from examining how the location of ports changed *within* cities over time. Figure F.2 shows that, over time (1953–2017), ports systematically moved *within city* to the outskirts, where land is typically less scarce. This came about as a combination of existing ports expanding outwards from the city center (by about one kilometer, on average), as well as new terminals being set up further from the city center (which were nine kilometers further form the city center, on average).

Taking the findings of this section together, we conclude that there is wide-ranging evidence for the increased land-intensity of containerization. This feature of the new technology mattered for where port development took place, both *across* and *within* cities.

4.3. 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:

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

Table 2

The local causal effect of shipping on population (Stylized fact 3).

Independent variables	Δln(Pop) (1)	Δln(Pop) (2)	Δ ln(Ship) (3)	Δln(Pop) (4)
Δln(Ship)	0.013 *0.052* (0.009)	0.006 *0.022* (0.073)		
Depth			0.272*** *0.134**** (0.086)	0.002 *0.003* (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. *** p<0.01, ** p<0.05, * p<0.1.

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 E.4 shows the panel specification allowing us to utilize the full decadal variation in the data.[13] 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.

These results are in contrast to 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 F.3), suggesting that North American cities may have had a larger than average population response to containerization.

We subject the 2SLS panel specification to the same set of robustness checks conducted above (Table E.5 and Appendix B.3). The coefficient is consistently small and indistinguishable from zero. 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 force. One potential channel may be the one working through the increased land-intensity of containerized ports. The results from Stylized fact 2 suggest these are empirically relevant and large enough to affect the economic geography of ports. In the next section, we build a model that incorporates this mechanism, thereby capturing both the benefits and the costs of port development.

5. A model of cities and endogenous port development

In this section, we present a flexible general equilibrium model that is consistent with the 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

---

[13] 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.

access, but also allows for two types of negative effects: the increased land use and the negative amenity externalities associated with port development.

5.1. 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.

5.1.1. 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.

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. This implies that the fraction of country-c workers who choose to live in city r of that country equals

$$\frac{N(r)}{\sum_{s\in c} N(s)} = \frac{\left[\frac{w(r)}{P(r)} a(r)\right]^{1/\eta}}{\sum_{s\in c}\left[\frac{w(s)}{P(s)} a(s)\right]^{1/\eta}} \tag{5}$$

where $w(r)$ is the nominal wage and $P(r)$ is the CES price index of consumption goods in the city.[14] Hence, $1/\eta$ equals the *migration elasticity*, the elasticity with which city population reacts to changes in the city's amenity-adjusted real wage. As $\eta \to 0$, an increase in the amenity-adjusted real wage attracts all workers to the city. This is because workers' idiosyncratic city tastes are identical and therefore play no role in their location decisions. At the other extreme, as $\eta \to \infty$, idiosyncratic city tastes completely dominate real wage differences, and each city is inhabited by those workers who prefer it over other cities for idiosyncratic reasons. Thus, an increase in the amenity-adjusted real wage leads to no change in the city's population.

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

$$a(r) = \bar{a}(r)\left[1 + Shipping(r)\right]^{-\rho} \tag{6}$$

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.

5.1.2. 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.[15] We normalize the stock of land available in each city to one.[16] Landlords have the same preferences over goods as workers. They do not work but finance 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. At the same time, more land for

---

[14] See Appendix C.2.1 for the derivation of Eq. (5).

[15] 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 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.

[16] 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 fixed in our model simulations.

transshipment necessarily implies less land available for production. In other words, the model implies a resource cost of land use that can influence the spatial allocation of port development across port cities, consistent with Stylized fact 2.

Port city landlords can charge a price for the transshipment service they provide. Competition among landlords drives down this price to marginal cost. Hence, profits from transshipment services are zero in equilibrium.[17]

5.1.3. 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.[18] The representative firm does not internalize the effect that its employment decision has on local population. Hence, it takes $N(r)$ as given.

5.1.4. 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[19]:

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

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,

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

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 $p_0, \ldots, p_M$ takes the form

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

where $\phi_{\varsigma}\left(d\left(r, p_0\right)\right)$ corresponds to the overland shipping cost between the origin and the first port en route p_0, and $\phi_{\varsigma}\left(d\left(p_M, s\right)\right)$ corresponds to the overland shipping cost between the last port en route p_M and the destination. $\phi_{\tau}\left(d\left(p_m, p_{m+1}\right)\right)$ denotes the sea shipping cost between ports p_m and p_{m+1}, a function of the minimum sea distance between the two ports, $d\left(p_m, p_{m+1}\right)$. Finally, $O\left(p_m\right)$ denotes the price that the firm needs to pay for transshipment services in port city p_m.[20]

---

[17] In Section 7.1, 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 C.7.

[18] 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).

[19] 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 (2023), finding optimal shipping flows is computationally more demanding.

[20] 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.

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

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

where $\nu\left(p_m\right)$ is an exogenous cost shifter capturing the fundamental efficiency of port p_m, $\psi\left(F\left(p_m\right)\right)$ is a non-negative, strictly decreasing and strictly convex function of $F\left(p_m\right)$, the share of land allocated to the port, and $Shipping\left(p_m\right)^{\lambda}$ captures congestion externalities arising from the fact that handling one unit of cargo becomes more costly as the total amount of shipping, $Shipping\left(p_m\right)$, increases for a given port size.[21] As each port city landlord is atomistic, she takes the price of transshipment services $O\left(p_m\right)$ and the total port-level shipping $Shipping\left(p_m\right)$ as given when choosing $F\left(p_m\right)$. Moreover, perfect competition among port city landlords ensures that the price of transshipment services is driven down to marginal cost and therefore

$$O\left(p_m\right)=\left[\nu\left(p_m\right)+\psi\left(F\left(p_m\right)\right)\right] Shipping\left(p_m\right)^{\lambda} \tag{7}$$

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 C.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.

5.1.5. *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 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 C.1 provides a formal definition and characterization of the equilibrium.

5.2. *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, $N(r)$:

$$N(r)^{[1+\eta\sigma+(1-\gamma-\alpha)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}}=\gamma^{\sigma-1}\tilde{a}(r)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} A(r)^{\frac{(\sigma-1)^2}{2\sigma-1}}(1-F(r))^{(1-\gamma)\frac{(\sigma-1)^2}{2\sigma-1}} MA(r) \tag{8}$$

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

$$MA(r)=\sum_{s=1}^{S}\frac{\tilde{a}(s)^{\frac{(\sigma-1)^2}{2\sigma-1}} A(s)^{\frac{\sigma(\sigma-1)}{2\sigma-1}}(1-F(s))^{(1-\gamma)\frac{\sigma(\sigma-1)}{2\sigma-1}} N(s)^{[1-\eta(\sigma-1)-(1-\gamma-\alpha)\sigma]\frac{\sigma-1}{2\sigma-1}}}{\mathbf{E}[T(r,s)]^{\sigma-1}} \tag{9}$$

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.

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. *Assume that Eqs. (6) and (8) hold. Then the population of city r is*

1. *increasing in the city's market access $MA(r)$ (market access effect);*
2. *decreasing in the share of land allocated to the port through its negative effect on $1-F(r)$ (land use effect);*
3. *decreasing in shipping flows through their negative effect on amenities $\tilde{a}(r)$ (disamenity effect).*

Proof. These results follow directly from Eqs. (6) and (8). ☐

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, thus increasing a city's market access. 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 presence of these opposing forces makes the model consistent with Stylized fact 3, that is, the fact that port cities' population remained constant in the data, despite the port development induced by containerization.

---

[21] To be precise, $Shipping\left(p_m\right)$ 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 $Shipping\left(p_m\right)$ as it simplifies the procedure of taking the model to the data.

Table 3
Taking the model to the data – summary.

Parameter	Description	Target
Step 1: Calculating shipping costs		
$t_\varsigma = 0.5636$	Overland cost elasticity w.r.t. distance	Allen and Arkolakis (2014)
$t_\tau = 0.0779$	Sea cost elasticity w.r.t. distance	Allen and Arkolakis (2014)
$\beta = 0.031$	Endogenous transshipment cost parameter	Correlation between shipping flows and port share in 1990
$B = 2.1$	Border trade cost	Ratio of international trade to world GDP in 1990
Step 2: Choosing the values of structural parameters		
$\alpha = 0.06$	Agglomeration externalities	Ciccone and Hall (1996)
$\gamma = 0.84$	Non-land share in production	Desmet and Rappaport (2017)
$\eta = 0.5$	Inverse of migration elasticity	Head and Mayer (2021)
$\sigma = 4$	Elasticity of substitution across tradables	Bernard et al. (2003)
$\rho = 0.005$	Disamenities	Port disamenities in Los Angeles (see Appendix C.9)
$\theta = 203$	Idiosyncratic shipping cost dispersion	Allen and Arkolakis (2019)
$\lambda = 0.074$	Congestion externalities in ports	Abe and Wilson (2009)
Step 3: Recovering post-containerization fundamentals		
$\bar{a}(r)$	Fundamental city amenities	Population, GDP and shipping by city in 1990
$A(r)$	Fundamental city productivity	
$v(r)$	Exogenous transshipment costs	

6. 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 seven structural parameters. Finally, we back out the values of unobserved city fundamentals that rationalize the post-containerization data.

6.1. Calculating shipping costs

To calculate the shipping costs across all potential routes, we need to specify 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

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

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).

Next, we specify 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 $\psi' < 0$.

Given that β drives the relationship between the value of shipping flows and the port share (see Eq. C.4 in Appendix C.2), we calibrate it to match the correlation between these two variables in the data. 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 F.4 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 (see Appendix D.5 for details).

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 $B = 2.1$.

6.2. 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 (1996). 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 (Table 3). Given we set $\alpha = 0.06$, this suggests choosing $\gamma = 0.84$. 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.

On the consumption side, we have three structural parameters. One is the parameter driving the migration elasticity, which we set to $\eta = 0.5$. The implied migration elasticity, $1/\eta = 2$, corresponds to the mean of the literature's existing within-country migration elasticity estimates, as reported in Table B.1 by Head and Mayer (2021). Another consumption-side parameter is the elasticity of substitution across tradable final goods, which we set to $\sigma = 4$ based on Bernard et al. (2003). Our final consumption-side parameter is 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); see section C.9 for details.

Finally, there are two structural parameters that influence the shipping technology. One is the dispersion 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 $\lambda = 0.074$.

6.3. 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 $V = 364$.[22]

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

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

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

Once population $N(r)$ and wages $w(r)$ are available for each city and the value of shipments, $Shipping(r)$, 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, $\tilde{a}(r)$, fundamental city productivities $A(r)$, and each port city's exogenous transshipment costs $\nu(r)$. We provide the details of this inversion procedure in Appendix C.3. The complex structure of the model does not allow us to prove that the inversion procedure identifies a unique set of $\tilde{a}(r), A(r)$ and $\nu(r)$. 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.

7. Counterfactuals

We conduct two counterfactuals in the model. The first counterfactual is backward-looking and 'rolls back' containerization. This allows us to estimate the aggregate effects of containerization. The second counterfactual is forward-looking and studies the port development undertaken by the Chinese government as part of their 'Belt and Road Initiative.' This allows us to illustrate the effects of targeted port development policy on targeted and untargeted cities, as well as aggregate welfare. Appendix C.4 discusses how we numerically solve for counterfactual equilibria in the model.

7.1. Rolling back containerization

In our first counterfactual, we compare the post-containerization (1990) equilibrium of the model to a counterfactual equilibrium in which containerization did not arise. When simulating the counterfactual, we account for the technological aspects of containerization in seaports that we document in Sections 2 and 4.1: lower costs, particularly in deep ports. This requires us to change the values of two model fundamentals relative to the post-containerization equilibrium.

---

[22] 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 $Ship(r)$ from zero to one in these cities.

First, 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 6.3, and $Depth(r)$ is the depth measure, defined in Section 4. 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 $\log \nu(r)$.

Second, 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 average change in the sum of exogenous *and* endogenous transshipment costs as a result of containerization.

We estimate that containerization reduced transshipment costs 25% by 1990 based on the following procedure. 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); see Appendix D.11 for details. Using the more conservative estimate of 70%, this gives us a 25% decrease in average transshipment costs. 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 $\nu_{CF} = 0.227$.

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 C.8 investigates how adding an additional inland cost increase to our counterfactual simulation changes the results.

The aggregate effects of containerization. We estimate that aggregate world welfare increased by 3.4% as a result of containerization. We define the change in aggregate world welfare as the average of changes in country-level expected welfare between the counterfactual and the 1990 equilibrium, weighted by country population. Within each country, labor mobility equalizes expected welfare across cities, as in Redding (2016). However, we do not allow for mobility across countries, hence different countries experience different welfare effects.

Consistent with Stylized fact 1, we find that containerization led to a boom in shipping flows. More precisely, containerization increased the international trade to world GDP ratio by 4.17 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 since the reduction in trade costs leads to increased demand for shipping, encouraging more investment in port development. Figure F.5 presents the full distribution of port share changes across cities. The median change is 2 percentage points, while the 5th percentile is −2 pp and the 95th percentile is 33 pp.

How large was the cost of increased land use due to containerization? To answer this question, 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 label the *resource costs* of containerization. Finally, containerization might yield gains from increased specialization of cities in port or non-port activities, which we label as the *specialization gains* from containerization.

To assess the quantitative importance of each of these margins, we develop two alternative models that we label as 'Benchmark 1' and 'Benchmark 2.' In 'Benchmark 1,' port development yields transshipment cost reductions but does not require land use. In 'Benchmark 2,' land needs to be used to reduce transshipment costs, but we restrict land use to be identical across port cities (and equal to the mean port share in our baseline). We provide a detailed description of each benchmark model and their quantitative estimation in Appendix C.5.

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 3.55% in Benchmark 1. In Benchmark 2, the gains from containerization reduce to 3.27%. The difference between Benchmark 1 and Benchmark 2, 0.28 percentage points, captures the resource costs of containerization. These costs are sizeable: they account for as much as 8% of the gains from the shipping cost reduction. Finally, the difference between Benchmark 2 and our baseline model, 0.13 percentage points, captures the specialization gains from containerization. Note that these gains are able to offset about 45% of the resource costs of containerization, but they do not fully compensate for all the costs.

In Appendix C.8, we show that the aggregate 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, different changes in exogenous transshipment costs, and a model in which landlords make profits from the provision of transshipment services.

Table 4
The effects of targeted port development: The Maritime Silk Road.

	Baseline				Benchmark 1			
	(1) Δ ln(Ship)	(2) Δ ln(Port cost)	(3) Δ ln(Market access)	(4) Δ ln(Population)	(5) Δ ln(Ship)	(6) Δ ln(Port cost)	(7) Δ ln(Market access)	(8) Δ ln(Population)
Treated port city	0.78406***	−0.14913***	0.02996***	−0.01237***	0.60662***	−0.10536***	0.02779***	0.00296**
	(0.11983)	(0.03097)	(0.00600)	(0.00336)	(0.09078)	0	(0.00623)	(0.00138)
Untreated port city in treated country	−0.41352***	0.01506**	0.01484***	0.00618***	−0.29749***	0	0.01674***	−0.00140
	(0.08223)	(0.00705)	(0.00325)	(0.00236)	(0.04999)	0	(0.00321)	(0.00125)
Port city in untreated country	0.01032	−0.00188	0.00040	−0.00072	0.00287***	0	0.00049***	0.00012***
	(0.01103)	(0.00785)	(0.00031)	(0.00053)	(0.00034)	0	(0.00003)	(0.00002)
Inland city in treated country			0.02247***	0.00128***			0.02214***	−0.00008
			(0.00134)	(0.00037)			(0.00133)	(0.00013)
Inland city in untreated country			0.00028***	0.00002			0.00029***	−0.00002***
			(0.00007)	(0.00002)			(0.00007)	(0)
Observations	553	544	2636	2636	553	553	2636	2636
R-squared	0.193	0.030	0.439	0.034	0.429	1.000	0.457	0.027

Notes: The regressors are dummy variables that divide the cities into 5 mutually exclusive groups as indicated, the regression is estimated without the constant. 'Treated port' indicates the 24 treated ports of the Maritime Silk Road counterfactual. 'Treated country' indicates countries that have at least one treated port. Robust standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

7.2. The effects of targeted port development

In our second counterfactual, 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.' 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. 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 F.6 for the set of targeted ports). We take the targeted ports from OECD (2018) and choose the decrease in exogenous transshipment costs, $\nu(r)$, to be 10% to illustrate the effects of a sizeable, but not dramatic decrease in costs. We keep all other fundamentals of the model fixed at their levels recovered in Section 6.

Table 4 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 7.1).

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 (column 1). 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 response to the policy leads to increasing returns from port development, drawing additional shipping into targeted cities and away from non-targeted ones. Table E.10 shows the results remain similar if we include country fixed effects.

We also study the effects on cities' market access (as defined by Eq. (9)) 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 Fig. 2. We find the most dramatic distributional effects in Asia. Strikingly, we see a dramatic reallocation of shipping to China and away from Singapore (which we estimate loses almost 50% of its shipping flows).[23] 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 increasing returns to port development. As shipping moves away from Singapore towards targeted ports, incentives to develop the port of Singapore decrease, which ultimately leads the city to cut back substantially on its port activities by reallocating land away from the port. However, 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 (at the 98th percentile of the world productivity distribution according to our model). 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

---

[23] It should 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.

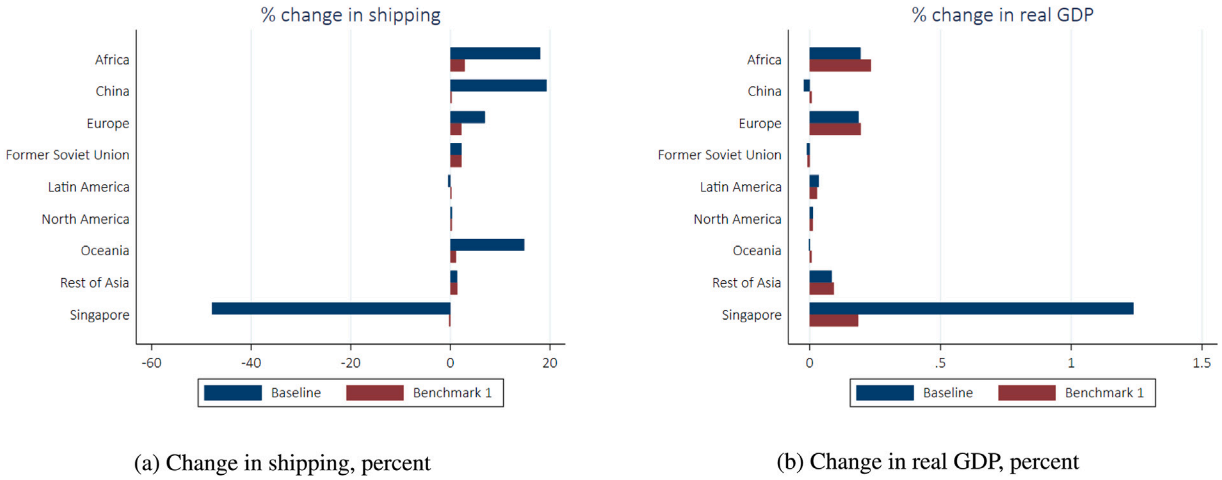

(a) Change in shipping, percent

(b) Change in real GDP, percent

Fig. 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 1990 equilibrium 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.

these additional factors. 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. 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.

Declaration of competing interest

The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

Data availability

Replication for Ducruet, Juhasz, Nagy and Steinwender "All aboard: The effects of port development (Original data) (Mendeley Data)

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jinteco.2024.103963.

References

Abe, K., Wilson, J., 2009. Weathering the storm: Investing in port infrastructure to lower trade costs in east Asia. World Bank Policy Research Working Paper 4911, World Bank.
Allcott, H., Keniston, D., 2017. Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. Rev. Econ. Stud. 85 (2), 695–731.
Allen, T., Arkolakis, C., 2014. Trade and the topography of the spatial economy. Q. J. Econ. 129 (3), 1085–1140.
Allen, T., Arkolakis, C., 2019. The welfare effects of transportation infrastructure improvements.
Allen, T., Atkin, D., 2022. Volatility and the gains from trade. Econometrica 90 (5), 2053–2092.
Altomonte, C., Colantone, I., Bonacorsi, L., 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. Amer. Econ. Rev. 69 (1), 106–116.
Armenter, R., Koren, M., Nagy, D., 2014. Bridges.

Bernard, A., Eaton, J., Jensen, J.B., Kortum, S., 2003. Plants and productivity in international trade. Am. Econ. Rev. 93 (4), 1268–1290.

Bernhofen, D., El-Sahli, Z., Kneller, R., 2016. Estimating the effects of the container revolution on world trade. J. Int. Econ. 98, 36–50.

Bleakley, H., Lin, J., 2012. Portage and path dependence. Q. J. Econ. 127, 587–644.

Brancaccio, G., Kalouptsidi, M., Papageorgiou, T., 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., Lin, J., 2022. Freeway revolts! the quality of life effects of highways. Rev. Econ. Stat. 1–45.

Brooks, L., Gendron-Carrier, N., Rua, G., 2021. The local impact of containerization. J. Urban Econ. 126 (C).

Ciccone, A., Hall, R., 1996. Productivity and the density of economic activity. Amer. Econ. Rev. 86 (1), 54–70.

Coşar, A.K., Fajgelbaum, P., 2016. Internal geography, international trade and regional specialization. Am. Econ. J. Microecon. 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., Neary, J., 1982. Booming sector and de-industrialisation in a small open economy. Econom. J. 92 (368), 825–848.

Coşar, A., Demir, B., 2018. Shipping inside the box: Containerization and trade. J. Int. Econ. 114, 331–345.

Desmet, K., Rappaport, J., 2017. The settlement of the United States, 1800–2000: The long transition towards Gibrat's law. J. Urban Econ. 98, 50–68.

Donaldson, D., Hornbeck, R., 2016. Railroads and American economic growth: A "market access" approach. Q. J. Econ. 131 (2), 799–858.

Donaldson, D., Storeygard, A., 2016. The view from above: Applications of satellite data in economics. J. Econ. Perspect. 30 (4), 171–198.

Ducruet, C., Cuyala, S., Hosni, A.E., 2018. Maritime networks as systems of cities: The long-term interdependencies between global shipping flows and urban development. J. Transp. Geogr. 66, 340–355.

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

Ducruet, C., Polo Martin, B., Sene, M.A., Lo Prete, M., Sun, L., Itoh, H., Pigné, Y., 2024. Ports and their influence on local air pollution and public health: A global analysis. Sci. Total Environ. 915, 170099.

Eaton, J., Kortum, S., 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., Redding, S., 2022. Trade, structural transformation and development: Evidence from Argentina, 1869-1914. J. Polit. Econ. 130 (5), 1249–1318.

Fajgelbaum, P.D., Schaal, E., 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. Econ. J. 86 (343), 536–550.

Fujita, M., Mori, T., 1996. The role of ports in the making of major cities: Self-agglomeration and Hub-effect. J. Dev. Econ. 49, 93–120.

Ganapati, S., Wong, W.F., Ziv, O., 2022. Entrepot: Hubs, scale, and trade costs. Am. Econ. J.: Macroeconomics (Forthcoming).

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., Mayer, T., 2014. Gravity equations: Workhorse, toolkit, and cookbook. In: Handbook of international economics, vol. 4, Elsevier.

Head, K., Mayer, T., 2021. The United States of Europe: A gravity model evaluation of the four freedoms. J. Econ. Perspect. 35 (2), 23–48.

Heiland, I., Moxnes, A., Ulltveit-Moe, K.H., Zi, Y., 2023. Trade from space: Shipping networks and the global implications of local shocks.

Holmes, T., Singer, E., 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. J. Econ. Perspect. 21 (3), 131–154.

Kahveci, E., 1999. Fast turnaround ships and their impact on crews.

Koenig, P., Pigné, Y., Poncet, S., Sanch-Maritan, M., Duvallet, C., 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. J. Dev. Econ. 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., Vallianatos, M., 2012. Importing Harm: U.S. Ports' Impacts on Health and Communities. Technical Report, Trade, Health and Environment Impact Project.

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

Moriconi-Ebrard, F., 1994. Geopolis.

Nagy, D.K., 2023. Hinterlands, city formation and growth: Evidence from the U.S. westward expansion. Rev. Econ. Stud. 90 (6), 3238–3281.

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.

Pines, D., Sadka, E., 1985. Zoning, first-best, second-best, and third-best criteria for allocating land for roads. J. Urban Econ. 17 (2), 161–173.

Port of San Francisco, 1971. San Francisco Port Needs, Shipping and Area Requirements. Technical Report, Port of San Francisco.

Redding, S., 2016. Goods trade, factor mobility and welfare. J. Int. Econ. 101, 148–167.

Redding, S., Turner, M., 2015. Transportation costs and the spatial organization of economic activity. In: Handbook of regional and urban economics, vol. 5, Elsevier, pp. 1339–1398.

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. Q. J. Econ. 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. Swedish J. Econ. 74 (1), 161–173.

Solow, R., Vickrey, W., 1971. Land use in a long narrow city. J. Econ. Theory 3 (4), 430–447.

The Economist, 2023. Investments in ports foretell the future of global commerce. In: The Economist. Available at: https://www.economist.com/interactive/business/2023/01/14/investments-in-ports-foretell-the-future-of-global-commerce. (Accessed: 22nd January 2024).

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. Am. Econ. J.: Appl. Econ. 14 (4), 127–166.