NBER WORKING PAPER SERIES

ALL ABOARD: THE EFFECTS OF PORT DEVELOPMENT

César Ducruet
Réka Juhász
Dávid Krisztián Nagy
Claudia Steinwender

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

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2020

We thank Treb Allen, David Atkin, Leah Brooks, Don Davis, Dave Donaldson, Joseph Doyle, Nicolas Gendron-Carrier, Matt Grant, Gordon Hanson, Tarek Hassan, Tom Holmes, David Hummels, Amit Khandelwal, Giampaolo Lecce, Nels Lind, Nina Pavcnik, Giacomo Ponzetto, Jim Rauch, Steve Redding, Roberto Rigobon, Andres Rodriguez-Clare, Esteban Rossi-Hansberg, Daniel Sturm, Tavneet Suri, Jaume Ventura, Jon Vogel and David Weinstein for helpful comments and discussions. We thank Bruce Blonigen for kindly sharing data and we thank staff at the Port Authorities of Houston, New Orleans, Long Beach, Los Angeles, Portland, San Francisco and Seattle for helping with our information requests. Olalekan Bello, Sabrina Chen, Naman Garg, Yi Jie Gwee, Hamza Husain, Felix Iglhaut, Rodrigo Martinez Mazza, Emanuela Migliaccio, Shuhua Si, Yue Yu, Howard Zihao Zhang and a team of Columbia University undergrad students provided outstanding research assistance. Reka acknowledges funding from the Provost's Office at Columbia. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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

ABSTRACT

This paper examines the effects of port development on the economy. By using scarce local land intensively, ports put pressure on local land prices and crowd out other forms of economic activity. We use the introduction of containerized shipping -- a technology that substantially increased land requirements at the port -- to estimate the effects of port development. We find an important role for the crowding-out effect both at the local and at the aggregate level. First, we show that the causal effect of the shipping boom caused by containerization on local population is zero -- port development increases city population by making a location more attractive for firms and consumers, but this well-known market access effect is fully offset by the crowding-out mechanism. Second, to measure the aggregate implications, we add endogenous port development to a standard quantitative model of cross-city trade. Through the lens of this model, we estimate that containerization increased aggregate world welfare by 3.95%. However, relative to the positive welfare effects of a trade-cost reduction in standard models, our model implies a sizeable welfare cost associated with the increased land-usage of ports, partly offset by welfare gains from endogenous specialization based on comparative advantage across port- and non-port activities. In terms of the distributional effects, we find that initially poorer countries gained more from containerization as they had a comparative advantage in port development.

César Ducruet
CNRS
13 rue du Four
Paris 75006
France
cdu@parisgeo.cnrs.fr

Réka Juhász
Department of Economics
Columbia University
420 West 118th Street
New York, NY 10027
and NBER
rj2446@columbia.edu

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

Claudia Steinwender
MIT Sloan School of Management
100 Main Street, E62-521
Cambridge, MA 02142
and NBER
csteinwe@mit.edu

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

Introduction

From Sri Lanka to the Netherlands, countries across the income distribution invest heavily in port development.[1] 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). Rich and poor countries alike view investments into ports as an integral part of their growth strategy, as modern facilities allowing for the fast flow of cargo through the port are a precondition for a country to participate in global production networks (Rodrigue, 2016, p. 131). Despite this, ports have been understudied relative to other forms of transport infrastructure such as roads or railways.[2] In particular, little is known about the economic effects of port development. What determines the economic geography of ports (i.e., where port activity is located)? What are the gains from port development and how are they distributed across countries?

In this paper, we study these questions by examining a breakthrough innovation in port technology: containerization, that is, the handling of cargo in standardized boxes. This new technology dramatically changed the transshipment of cargo at seaports during the 1960s and 1970s. 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 land-intensity of ports. Different to other transport infrastructures such as railways or roads, ports are investments that occupy large amounts of land in the cities in which they are located. For example, the ports of Antwerpen and Rotterdam occupy more than 30% of the metropolitan area of the city (OECD, 2014). By using locally scarce land resources heavily, ports drive up land rents and *crowd out* other economic activity.

In our analysis, we exploit the introduction of containerization to identify the crowding-out effect of port development. Using rich historical evidence, as well as detailed time series data on wharf dimensions for one port (New Orleans), we show that containerization is a much more land-intensive technology than the one it replaced. That is, more land is needed at the port under containerized technology; the data from New Orleans suggest that the land intensity of transshipment technology increased by about 75% after containerized terminals were introduced. The benefit of containerization was that by increasing land usage at the port, transshipment times (Port of San Francisco, 1971) and costs (Hummels, 2007) were drastically reduced.

We use a unique dataset of city populations and shipping flows worldwide for the period 1950-1990 to estimate the local, city-level effects of containerization. To isolate exogenous variation,

[1]As an example, the Port of Rotterdam (Netherlands) undertook the expansion of its container facilities by 110 ha in 2004 at a cost of EUR 657m, 200m of which was financed by the European Investment Bank (Source: https://www.eib.org/en/projects/pipelines/all/20030288). The Port of Colombo (Sri Lanka) has made massive investments in recent years. A single project upgrading harbor infrastructure was undertaken between 2008-2012 at a cost of Rs 42 billion (Source: https://www.slpa.lk/port-colombo/projects).

[2]Redding and Turner (2015) provide an overview of this literature. An exception is Brooks, Gendron-Carrier, and Rua (2019), who study the reduced-form effects of containerization on county-level economic outcomes in the U.S.

we build on a previous literature that has shown that access to deep sea ports was an important determinant of a city's suitability for containerization (Brooks et al., 2019; Altomonte, Colantone, and Bonacorsi, 2018). We develop a novel measure of 'naturally endowed' depth (as distinct from depth attained by dredging) using granular data on oceanic depths around each city in our data. We show that cities exogenously more suited to containerization witnessed a boom in shipping flows after the onset of containerization, but not before. Surprisingly, however, this boom in local shipping *did not* translate into population inflows: we find an effect of shipping on population in our IV estimates that is both economically and statistically insignificant.

We view the zero local population effects of containerization as an unexpected finding. It is in contrast to standard models that predict an inflow of population as improved market access makes a location more desirable for firms and consumers (Coşar and Fajgelbaum, 2016; Nagy, 2018; Fajgelbaum and Redding, 2018). Indeed, other papers studying similar shocks to a location's accessibility have found a positive effect on population (Bleakley and Lin, 2012; Campante and Yanagizawa-Drott, 2018; Brooks et al., 2019). However, the higher land intensity of containerized port technology can provide an explanation for the zero population effect. Intuitively, the increased use of scarce local land can counteract the market access effect by driving up land prices and crowding out other economic activity from the city. Consistent with an important role for land prices in determining where port development takes place, we indeed show that shipping increased disproportionately more in low land rent cities.

Informed by the local, reduced form effects of port development, in the second part of the paper we develop a tool for quantitative general equilibrium analysis. The model is an otherwise standard economic geography model of trading cities to which we add an endogenous port development decision. As such, the model incorporates not only the standard market access effect, but also allows for port development to crowd out other forms of economic activity. This is because in the model, developing the port (and hence reducing trade costs) requires scarce local land that can be used for other purposes. Whether a city ultimately gains in population is the outcome of the trade-off between the market access and crowding-out mechanisms. Thus, the model has the ability to rationalize the zero population effects of shipping found in the data.

Guided by the model, we re-estimate the causal effect of increased shipping flows on population controlling for market access. In line with the predictions of the model, our causal estimates point to a *negative* effect of shipping on city population once market access is controlled for. This finding provides further empirical evidence consistent with the crowding-out effect of port development.

In the final part of the paper, we quantify the aggregate and country-level effects of containerization 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. Next, we simulate the pre-containerization equilibrium in the model by *undoing* the transshipment cost reduction and increase in the land intensity of port technology that containerization caused. Comparing these two equilibria reveals the effects of containerization. We test whether the model can replicate the same local effects of containerization that we found in the reduced form. First, we show that the model-simulated data closely matches the zero population effects of shipping using the same IV strategy (based on depth) as in the reduced form. Second, we show that containerization increased shipping more in low land-rent cities, as in the data.

Our results show that containerization increased world welfare by 3.95%. To better understand how the crowding-out channel affects these welfare gains, we compare the aggregate welfare effects in our model to what a standard model in which transport cost reductions are *exogenous* and *free* (i.e., they do not use scarce resources) would predict. We find a quantitatively meaningful role for two mechanisms. First, we estimate the aggregate resource cost of containerization to be substantial: it offsets about 13% of the welfare gains arising from a standard model. Second, we also find a role for additional welfare gains stemming from endogenous specialization in port- and non-port activities based on comparative advantage. In particular, these gains offset about 63% of the resource cost of containerization. In addition, we find that, unlike in our model, the local population effects of shipping are positive, economically meaningful and statistically significant in the standard model. This result again underscores the link between the zero local population effects of shipping and the endogenous crowding-out mechanism that is present in our model.

Finally, we examine the distributional implications of containerization by studying the country-level welfare gains from containerization implied by the model. These gains are heterogeneous: while 26% of countries experience gains below 2%, 29% of them see gains above 10%. We find that initially poorer countries gained more from containerization. We show that in the model, this relationship is explained by poorer countries being less productive in non-port activities (hence having a comparative advantage in port activities) and these countries having worse market access before containerization. We also show that the negative relationship between pre-containerization market access and the gains from containerization is amplified relative to a standard model without endogenous port development. This is because endogenous port development can be conducted at a lower cost in poor countries that tend to have lower land rents on average. These findings highlight the importance of accounting for the endogenous crowding-out mechanism when quantifying how the gains from containerization were distributed across countries. They also imply that the 'Sri Lankas' of the world have more to gain from port development than the 'Netherlands'.

Related literature. A recent, growing literature provides evidence that better trading opportunities lead to local benefits inducing city development (Bleakley and Lin, 2012; Armenter, Koren, and Nagy, 2014; Nagy, 2020; Campante and Yanagizawa-Drott, 2018). Some of these studies focus on city development at port locations in particular (Fujita and Mori, 1996; Coşar and Fajgelbaum, 2016; Fajgelbaum and Redding, 2018). We contribute to this literature by showing that trade-induced development can also have substantial local costs. The crowding-out mechanism that drives the cost side in our setting also relates the paper to the 'Dutch disease' literature. This literature shows that booming industries can entail significant costs by putting a strain on scarce local resources and therefore crowding out other (tradable) sectors (Corden and Neary, 1982; Krugman, 1987; Allcott and Keniston, 2017).[3] 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 these two, seemingly disparate 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 trade across multiple countries with various dimensions of heterogeneity (Anderson, 1979; Eaton and Kortum, 2002; Melitz, 2003). These seminal models characterize trade and the distribution of economic activity across countries 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 (2020), who consider endogenous road construction in multi-location models of economic geography, as well as Brancaccio, Kalouptsidi, and Papageorgiou (2020), who endogenize trade costs in the non-containerized shipping sector. Unlike these papers, we focus on port development as a source of endogenous shipping costs, and solve for the decentralized equilibrium as opposed to the optimal allocation to quantify the effect of containerization-induced 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.[4] In particular, there is a growing empirical literature studying the effects of containerization (Hummels, 2007; Bernhofen, El-Sahli, and Kneller, 2016; Gomtsyan, 2016; Coşar and Demir, 2018; Holmes and Singer, 2018; Altomonte et al., 2018; Brooks et al., 2019) or the role of container shipping networks in world trade (Wong, 2017; Heiland, Moxnes, Ulltveit-Moe, and Zi, 2019; Ganapati, Wong, and Ziv, 2020). Most closely related is Brooks et al. (2019), 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, motivated by the evidence that containerization dramatically increased land use in ports, this paper highlights the crowding-out effect of containerization and finds sizeable local and global costs stemming from this effect. Sec-

[3] Another related paper is Falvey (1976), who discusses how the transportation sector can draw away resources from tradables in particular.

[4] Redding and Turner (2015) provides an overview of recent developments in this literature.

ond, to the best of our knowledge, this 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 main features of containerized technology. Section 2 discusses the main data sources used in the analysis. Section 3 presents the reduced form empirical strategy and results, while Section 4 introduces the model. Section 5 revisits the empirics guided by the predictions of the model. In Section 6, we take the model to the data, while in Section 7 we present our estimates of the aggregate effects of containerization. Section 8 concludes.

1 Containerization and other new port technologies

The introduction of steamships and railroads in the 19th century substantially reduced both water and overland transportation costs. However, transshipment technology – that is, the loading and unloading of cargo at transportation nodes such as seaports – remained slow and expensive (Krugman, 2011). As a report by McKinsey highlighted; "The bottleneck in freight transport has always been the interface between transport modes, especially the crucial land/sea interface" (1972, pp. 1-3). Containerization, that is, the handling of cargo in standardized boxes, was the breakthrough innovation that dramatically changed transshipment technology and reduced costs (Hummels, 2007; Rodrigue, 2016).[5] Within the space of a few years in the 1960s and 1970s, the technology ports used to transship cargo changed dramatically.[6]

In this section, we discuss features of containerized technology important for our analysis. First, we show that substantial transshipment cost reductions were achieved in shipping as a result of containerization. However, this came at the cost of needing to dedicate much more land to the port. Second, we show that many of the changes that we refer to using the shorthand term 'containerization' also affected non-containerized cargo. This point is important, as both our empirical and structural analysis capture the effects of new technologies on *all* cargo types.

1.1 The cost – space trade-off in containerization

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. The reason for this was that

[5]The key advantage of containerization is the speed and efficiency with which containers can be transshipped between different modes of transportation such as across ships, trucks and railroad (i.e., their intermodality), and within the same mode (e.g., ship-to-ship or truck-to-truck). Our objective in this paper is to understand the effect that containerization had on transshipment at *seaports*. This includes transshipping from one ship to another, as well as transshipping between vessels and railcars or trucks at the seaport. However, we do not analyze the effects containerization has had on inland transshipment, e.g., from road to rail.

[6]Rua (2014) provides a discussion of the swift adoption of containerization worldwide.

cargo came in many different sizes and so needed to be handled individually.[7] We illustrate breakbulk technology at work in the Port of New Orleans in 1954 in Panel A of Appendix Figure A.1 (the third-largest U.S. port in 1950 according to our data).[8] The San Francisco Port Commission (1971) estimated that it took 7-10 days to merely discharge cargo from a ship using this technology. According to Bernhofen et al. (2016), two-thirds of a ships' time would be 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.[9]

Port technology changed dramatically starting in the late 1950s when U.S. shippers first started placing cargo into boxes called containers.[10] Containerized port technology can be seen in its mature form at the Port of Seattle in 1969 (the seventh-largest U.S. port in 1970 according to our data) in Panel B of Appendix Figure A.1 (a mere 15 years after the photo at the Port of New Orleans was 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.

Cargo packed into containers at the origin and not opened until the final destination substantially reduced transshipment costs for a number of reasons.[11] First, as containers could be handled in a uniform way, loading and unloading times were vastly reduced. The San Francisco Port Commission (1971) estimated that a container ship could be unloaded and loaded in 48 hours or less at that time, a tenth of the previous time spent in port. Similarly, using detailed data on vessel turnaround times for one anonymized port, Kahveci (1999) estimates that the average time ships spent in port fell from 8 days to 11 hours as a result of containerization, a reduction of 94%. Second, the reduction in turnaround time justified investment into much larger vessels (Gilman, 1983). The average size of newly-built container ships increased by 402% between 1960 and 1990, as Appendix Figure A.2 shows. Larger ship sizes made it possible to realize even larger cost reductions through economies of 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

[7] By the 1950s, machinery was widely used across ports in the form of forklifts, conveyor belts and small cranes (Levinson, 2010, p. 18), but they did not eliminate the need to handle cargo individually, which was the main driver of lengthy transshipment times. The usage of machinery in breakbulk shipping is nicely illustrated in Figure A.1 that shows a small crane being used to receive cargo.

[8] The wharf shown in the figure was a newly completed extension to the import-export facilities of the port in 1954, suggesting that this was considered state-of-the-art technology as late as the mid-1950s. This is also evidenced by the use of cranes to offload cargo.

[9] 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).

[10] Containerized shipping was initially introduced on domestic routes between U.S. ports, but the technology was rapidly adopted and importantly, standardized worldwide in 1967 (Rua, 2014).

[11] In the following, we discuss transshipment cost reductions at ports, which is the focus of our paper. A more detailed discussion about other transport cost reductions as a consequence of containerization can be found in Rodrigue (2016).

50%.[12, 13]

Adapting ports to containerized technology was not without costs, however. Most importantly, faster turnaround times could only be achieved at the cost of building much larger terminals. This is a well-known feature of containerized ports in the transportation literature: In discussing the 'challenges' associated with containerization, Rodrigue (2016, p. 118) puts site constraints in the first place, and in particular, the large consumption of terminal space. Containerized terminals need more space as it is the easy accessibility of the containers that allows for efficient on- and off-loading. The containers are lined up next to where the ships dock, and space is also needed to rapidly off-load cargo. There are additional dedicated 'upland areas' near the facility that allow for the containers to be temporarily stored (New York Port Authority, 1958, p.5) and new space needed to be made for large 'railyards' where containers could await transshipment onto rail carriages (Riffenburgh, 2012, pp. xi-xii).[14]

The increased space requirements of containerized facilities were evident from the earliest days of the new technology.[15] As early as 1958 (two years after the first containerized shipments had sailed from New York), the New York Port Authority put in place plans to develop the Elizabeth facility for containerized cargo handling; "Extensive supporting upland area is one of the most important features of the development, since these large open spaces are indispensable in the handling of general cargo in the age of container ships" (1958, p. 5). The Port of San Francisco (the fifth largest port in the U.S. in 1950 according to our data) was raising alarm bells about the inadequacy of the city's finger piers to accommodate new types of cargo handling; "The Port [should] commence the phasing out of finger piers. [The piers are] commercially obsolete for the new generation of ships and the new types of cargo handling technology" (Port of San Francisco, 1971, p. 27). "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 in densely built up areas such as Manhattan and San Francisco were almost certainly doomed to decline as one observer noted for San Francisco; "Rows of finger piers adjacent to a

[12] The 6,000 TEU landmark for vessel size was surpassed in 1996, after our sample period (Rodrigue, 2016, p. 118).

[13] As more cranes can be used to unload the cargo of larger vessels, transshipment times did not need to increase substantially.

[14] Of course, warehouses and transit sheds were replaced to a large extent as containerization was rolled out. However, the space requirements of the two are not the same, as warehouses and transit sheds tended to be multi-story.

[15] Before containerization, there were some smaller innovations in handling breakbulk cargo such as palletization (whereby goods are placed on a pallet and handled as a unit) and pre-slinging (whereby goods are grouped together using slings and the unit is handled together). UNCTAD (1971) gives a more detailed overview of these trends. These new processes also allowed cargo to be handled as a unit, saving some transshipment time, though importantly, *not* in a standardized way (given the different-sized cargo involved). In Appendix A we show that these smaller improvements made contemporaries aware of the space – cost trade-off, and small changes to port layout (such as the widening of finger piers) were implemented. However, these pre-containerization changes were minor compared to the effect of containerization.

densely built up city could not adequately serve container shipping, which involved larger ships that required larger wharves and much larger areas of open space for loading and unloading" (Corbett, 2010, p. 164).

While it is difficult to quantify precisely how much more space containerized ports require, we have been able to find high quality data for one port that allows us to give a tentative answer: New Orleans. Appendix Figure A.3 presents data on the wharf length and area dimensions of terminals at the Port of New Orleans for the years 1950-1985.[16] The measure we are interested in is the area required to serve a ship under the different technologies. In addition, as the discussion above has highlighted, containerized ships were typically larger and required larger wharves, thus rendering a raw comparison between containerized and non-containerized terminals biased. For this reason, we examine the area of the port divided by wharf frontage (that is, the length of the wharf where ships can dock, thus accounting for differences in ship size). The numbers are striking; after containerization was introduced at the Port of New Orleans, the area per wharf frontage increased by 75%. Taking the rich historical evidence together with the data from New Orleans, we conclude that increased land intensity is an important feature of containerized technology.

1.2 New technologies for other types of cargo handling

To what extent did other types of non-containerized cargo handling benefit from similar innovations, or from spillovers from containerization? This is an important question, as our empirical research design will estimate the effects of new methods of cargo handling on the *sum* of shipping flows (that is, on the sum of containerized and non-containerized ships). The evidence suggests that non-containerized cargo transshipment also underwent some similar – albeit generally weaker – trends during our sample period. First, we see substantial reductions in ship turnaround times. Using data for one anonymized port between 1970-1998, Kahveci (1999) shows that ship turnaround times decreased by between 40-94% for different types of cargo, with the smallest gain being achieved for petroleum products and the largest for breakbulk cargo that was replaced by containers. Cars (for which the 'roll-on roll-off' technology was introduced; a technology that most closely resembles containerization in many ways) had almost the same efficiency gains as breakbulk cargo, followed by forest products and liquid bulk. Accordingly, faster turnaround times made larger ships cost-efficient, so vessel sizes also became significantly larger across many different cargo types. As Appendix Figure A.2 shows, the average size of newly built non-container vessels increased by 68% between 1960 and 1990 (in contrast to the 402% increase for container ships discussed above). It is also conceivable that non-containerized cargo benefited from other spillovers from the containerized technology. For example, deeper ports made it possible to receive

[16] Detailed data on wharf dimensions are reported in various editions of the port's annual reports (see the notes to Figure A.3 for information on sources). These data include area information on warehouses that may serve multiple terminals, even if they are not located directly at the terminal.

larger ships for all types of cargo.

For these reasons, we will use the term 'containerization' as shorthand for the bundle of new technologies and possible spillover effects of containerization that affected all cargo.

2 Data

Our analysis builds on a decadal city-level dataset of shipping flows, population, and other economic outcomes for the period 1950-1990. We complement this with GIS data that allows us to calculate certain geographic characteristics of the city important for our analysis. We review the main variables used in the analysis below. Summary statistics for the main variables are reported in Appendix Table A.1. Additional details on data construction and data sources are discussed in Appendix C.

Shipping Flows. Crucial to our analysis is a dataset of worldwide bilateral ship movements at the port level. These data correspond to the period 1950-1990, and come from Ducruet, Cuyala, and Hosni (2018). An observation is a ship moving from one port to another at a particular point in time. As such, it is similar to contemporary satellite AIS (Automatic Identification System) data that tracks the precise movements of vessels around the globe.[17] One week samples of these data were extracted from the *Lloyd's Shipping Index*, a unique source that provides a daily list of merchant vessels and their latest inter-port movements.[18] We are aware of no previous application in the economics literature.

These data provide us with rich variation to study the geography of sea-borne trade through the second half of the 20th century. First, they cover both domestic and international shipping. Second, the data include both containerized and non-containerized cargo. Third, 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*. We clean these data by hand-matching them to the 1953 and 2017 editions of the *World Port Index (WPI)*, which is a widely used reference list of worldwide ports.[19] Our base sample consists of

[17]These type of AIS data are used in Heiland et al. (2019) and Brancaccio et al. (2020).

[18]The data were entered from issues of the *Lloyd's List* for the first week of May. It should be noted that ship movements from the first week of May will often include journeys that took place in March or April due to the time lag between sailing and printing. The data are discussed in more detail in Ducruet et al. (2018).

[19]The initial Lloyd's List sample of 'ports' included ports on navigable rivers such as Budapest, Hungary. We therefore chose to discipline the sample of ports using WPI. We use a historic and current edition of the WPI to ensure we capture both ports that may no longer exist, and ones that only appear later in the period. A different approach would have been to choose a distance threshold from the coast and drop any port located further from the coast than the threshold. This definition, however, is very sensitive to the precision of the coastline shapefile used to calculate

Lloyd's List ports that match to at least one of the WPI editions.

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. For example, New York (New York) and Newark (New Jersey) form one 'city' according to this definition. As is common with censored city population data, 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. This type of sample selection can lead to important biases as we oversample fast-growing cities that were initially small. In the empirical section, we 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. Appendix Table A.2 contains the breakdown of port and non-port cities. Of the 2,636 cities in the Geopolis dataset, 553 have at least one port assigned. We label these as *port cities*. For these cities, we observe shipping flows and city population for the years 1950-1990. In addition, we have information for 1,592 ports that are not assigned to a city in the Geopolis dataset. Appendix Figure A.4 visualizes the spatial distribution of cities in our data, distinguishing between port and inland cities. Our reduced form empirical analysis focuses on estimating the local effects of containerization on the 553 port cities. The quantitative estimation covers the full set of 2,636 Geopolis cities (port and non-port cities).

Underwater elevation levels. 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 the city.[20]

Saiz land rent proxy. While we are not aware of any dataset that covers land rents globally going back to the 1950s, Saiz (2010) has proposed a geography-based measure that correlates well with

distance form the coast, which is why we did not choose this method. Despite filtering the Lloyd's List sample through the WPI, our final sample still contains a handful of ports that are very far inland. In the empirical analysis, we show that our results are robust to different ways of treating these 'inland ports'.

[20]Additional information on these data are described in Appendix C.1.

land-rents. This allows us to construct land rent proxies for all cities in our dataset. The 'Saiz-measure' is defined as follows: Take a 50 kilometer radius around the centroid of the city. Exclude all sea cells, all internal water bodies and wetland areas and all cells with a gradient above 15%. The remaining cells, as a share of the total cells can be used as a proxy for land rents. We replicate the methodology in Saiz exactly, using GIS data that have global coverage.[21] Spatial variation in the Saiz measure is visualized in Appendix Figure A.5.

City-level GDP per capita. 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).[22]

Port share. We define the port share as the share of a city's land occupied by the port. These data are needed for calibrating the model for the quantitative estimation. We have been able to find high-quality, consistent data for the land area occupied by ports for only 7 port cities in 1990, as these data are typically not recorded.[23]

3 The reduced form effects of containerization

In this section, we study the local effects of containerization on port cities. To isolate the causal effect of containerization, we develop an exogenous measure of port suitability based on the depth of the sea around the port. We examine three questions. First, did cities exogenously more suited to containerization witness an increase in shipping flows? We confirm that they did, but only after 1960, consistent with historical evidence on timing. Second, did this boom in shipping flows translate into increased city population? Surprisingly, we find no discernible causal effect of shipping flows on local port city population. Third, we work towards understanding this result by examining whether our data show evidence consistent with the land-intensive nature of shipping discussed in Section 1. We find that indeed, low-rent cities witnessed higher increases in shipping flows, as the land price mechanism would suggest. We begin this section by introducing the exogenous measure of port suitability used throughout the paper and proceed to discussing the three empirical results.

[21] These sources are documented in Appendix C.2.

[22] More details on this exercise are provided in Appendix C.3.

[23] We provide additional details on data construction in C.4.

3.1 An exogenous measure of port suitability

Section 1 discussed the fact that containerization led to larger ship sizes, and that 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., 2019; 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 certain buffer rings around the geocode of the port and sum the number of cells that are 'very deep', which we define as depth greater than 30 feet following Brooks et al. (2019). The authors argue that given vessel sizes in the 1950s (pre-containerization), depth beyond 30 feet conferred no advantage to the port.[24] Below, we will test how reasonable this assumption is by examining pre-trends in shipping.

To operationalize our measure, we need to take a stand on which set of cells around the port to consider. Our aim is to measure depth in areas around the port that are used by ships to navigate and wait for their docking time. To get a sense of where these areas are for a typical port, we examine the location (using exact geocodes) of stationary ships around the port in a one hour window for 100 random ports in our sample using contemporary data.[25] The cumulative distribution of ships around these ports, split by percentiles of port size, is shown in Appendix Table A.3. For these 100 ports, we find stationary ships located up to 25-30 km around the port, and sometimes even beyond.[26] However, the majority of stationary ships are located within 5 km, which justifies our baseline measure of port suitability: the log of the sum of 'very deep' cells in a buffer ring 3-5 km around the port.[27] We will examine the effect of depth measured at various buffers and show that the effects are similar in nearby rings, suggesting that the variation we use from the 3-5 km buffer is a representative measure of depth at the port.

Testing for endogenous dredging. 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,

[24] In practice, the geocode of the port is typically not exactly on the coastline. To correct for the measurement error that we would introduce from having geocodes closer or farther away from the coastline, we project all geocodes onto the coastline.

[25] These data are from *marinetraffic.com* and refer to *stationary* ships near the port captured between November 4 and 10, 2019, at 12:00-13:00 local time. More details regarding these data are provided in Appendix C.8. There is a concern that measures of where ships are found around the port *today* is a poor proxy for where ships were located during our sample period. Partly for this reason, we will show that depth measured in the same way at different nearby buffers yields similar results.

[26] We observe stationary ships farther away from the port for larger ports – in particular those in the 75th-100th percentile, which makes sense given that larger ports need to accommodate more ships at any given time. We would therefore also expect there to be more ships waiting around the port for their docking time.

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

they typically do not dredge entire areas in our buffers, but narrow channels that ships use to navigate to the port. By calculating depth over many sea cells, the vast majority of depth measurements for each port should reflect naturally endowed depth. We test this assumption in the following way. For 100 random ports in our sample, we obtained access to nautical maps from *marinetraffic.com* which clearly demarcate the dredged channels that ships use to navigate to the port.[28] We then constructed a binary variable, *'Dredging'*, that takes the value 1 if a port has a dredged channel in the 3-5 km buffer ring used in the baseline. Appendix Table A.4 shows the association between this measure and the depth measure. The unconditional association (column 1) is *negative* and statistically significant. That is, ports that we measure to be shallow are more likely to have a dredged channel. This is what we would expect to find if our measure captured naturally endowed depth.[29]

Balancing checks. Finally, we examine the extent to which our measure of exogenous port suitability is correlated with other observables pre-containerization in order to assess the types of confounders that may bias the results. Appendix Table A.5 shows the results. If greater depth would have led to more shipping even before containerization, we would expect to see a positive coefficient between depth and shipping flows. However, we see that the unconditional measure of depth is *negatively* correlated with both the level of shipping flows in 1950 (measured in logs), and population in 1950 (measured in logs), indicating that initially small cities had larger depth. In terms of growth rates pre-containerization, depth is weakly positively correlated with population growth between 1950 and 1960 (the coefficient is significant at 10%). This suggests that our depth measure is correlated with small cities that are growing relatively fast, i.e., population convergence. In order to purge our depth measure of this variation, we residualize it on city population in 1950 (measured in logs).[30] We re-examine how the part of the variation in depth that is uncorrelated with 1950 population, 'residualized depth', correlates with the same observables. Reassuringly, residualized depth is correlated neither with the level of shipping and population in 1950 (the latter by construction), nor with the change in shipping and population between 1950 and 1960. In the empirical analysis, we therefore use the residualized measure of depth as the baseline measure of exogenous port suitability.

Appendix Table A.5 also shows the correlation with other observables. Residualized depth is uncorrelated with country level GDP per capita measured pre-containerization and the latitude and longitude of the city (this is also true for raw depth). However, both depth and residualized

[28] For more details on this exercise, see Appendix C.9.

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

[30] More precisely, we regress the log of depth on the log of population in 1950 and take the residuals from this regression. Population in 1950 is not observed for 21 out of 553 port cities. For these, we replace 1950 population with the first year in which population is observed, which is generally 1960.

depth are correlated with the Saiz land rent proxy. This is perhaps unsurprising, as arguably similar geographic characteristics determine the overland (Saiz measure) and underwater (depth measure) geographic features around a city. For this reason, we show robustness of all our results to the inclusion of the Saiz land rent proxy interacted with year indicator variables. Appendix Figure A.6 visualizes the spatial variation in the residualized depth measure. While there seems to be some spatial correlation in depth, there is also a fair amount of variation within narrowly defined regions. We will tackle the issue of spatial correlation head-on in the empirics by testing the robustness of our results to using only within-region variation. We will also show that our results are robust to adjusting for spatial autocorrelation in the error term by reporting Conley standard errors (Conley, 1999).

3.2 Results

In this subsection, we use the depth-based measure of port-suitability to examine the local effects of containerization.

Result 1: Depth predicts shipping, but only after 1960. First, we examine whether depth predicts shipping flows during our sample period. We implement this using the following flexible specification that allows us to examine the timing of when depth started to matter for shipping.

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

The outcome variable of interest, $\ln(Ship_{it})$, is the log of shipping flows observed in city i at time t. This is the sum of all shipping flows recorded for city i at time t. We expect containerized technology (as defined in Section 1) to affect this measure both through its effect on containerized and non-containerized cargo. We need to take a stand on the treatment of zeros in the shipping data.[31] In the baseline measure, we annualize the weekly counts of ships from the raw data by multiplying the one-week sample of shipping flows we observe by 52. This is primarily so that our results are comparable to regressions we run using model-simulated data in the quantification exercise in Section 7. Finally, we replace the zeros in the data with ones and take the natural logarithm of this adjusted annualized count.[32] $Depth_i$ is the cross-sectional measure of port suit-

[31] The data contain zeros for two reasons: First, we may observe zeros because of measurement error: small ports with low shipping flows may not register an interport-movement during the week in which we capture the data. Second, zeros may appear due to the time-invariant definition of port status that we use. We observe zero shipping flows in a particular year if a port was established in the city only after 1950, or if a port shut down in the city during our sample period. Overall, we observe zero shipping flows for 16% of the port-year observations. From examining the data, the zeros seem to be more likely driven by mismeasuring small shipping flows rather than the entry and exit of ports.

[32] In robustness checks discussed below, we show that all of the results presented in this section are robust to other

ability defined in the previous subsection. We interact this measure with binary indicators for the decades 1960 – 1990 to estimate the time path of how depth affected shipping flows. In addition, we include the full set of city and year fixed-effects (denoted α_i and δ_t, respectively) as well as the log of population in 1950 interacted with year indicator variables across all specifications. This is equivalent to using the residualized depth measure in a panel setting. We cluster standard errors at the city level in the baseline to account for the serial correlation of shocks. We also report Conley standard errors (in curly brackets). As these are always very close to the clustered standard errors and do not change the statistical significance of our results, we will only report them for the main results.[33] 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. A number of points should be noted. First, deeper ports did not witness differential growth in shipping flows between 1950 and 1960 (coefficient -0.051, se. 0.063), consistent with this being a decade in which containerization was just being developed in a few ports around the world. Second, we see an effect of depth in each of the following decades, as containerization was adopted worldwide. The coefficient of interest is much larger and significantly different from zero for the interaction of depth and each year indicator including and after 1970 (e.g., the coefficient for the 1970 interaction is 0.222, se. 0.069). This is consistent with containerization technology being rolled out in the early 1960s across US ports and worldwide later in the decade, as we discussed in Section 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 some evidence that this assumption is plausible. Next, we turn to further testing this result with more demanding specifications. One concern is that many determinants of depth may be spatially correlated and if true, the estimates could be hard to disentangle from broader regional trends. To this end, column (2) adds the full set of 'coastline' by year-fixed effects to examine the extent to which our identifying variation relies on cross-regional variation.[34] Note that this set of fixed effects subsumes continent by year fixed effects. A comparison of the coefficients between columns (1) and (2) reveals that they are very similar, suggesting that broader regional trends are unlikely to be driving the effects.

Column (3) adds the Saiz land rent proxy interacted with year indicators to capture trends

standard ways of dealing with the zeros. In these, we do not annualize the data in order to verify that this transformation does not drive the results.

[33] We allow for spatial correlation at distances up to 1,000 km and set the spatial decay function to be linear.

[34] 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'. Appendix Figure A.7 visualizes the 'coastlines' that this exercise yields.

driven by the time-varying effect of land rents. Recall that this is a particularly important robustness check as the depth measure is correlated with the Saiz measure. Column (3) shows that the results are robust – the coefficients of interest become a bit larger across the board but the pre-trends remain small and statistically indistinguishable from zero, while the estimated coefficients including and after 1970 are much larger and highly significant throughout. Column (4) adds country GDP per capita (measured in 1960) interacted with year indicators to control for potentially differential growth trends across initially rich and poor countries.[35] The coefficients are remarkably stable.

Based on these results, we introduce a 'containerization' treatment indicator that turns on in years including and after 1970. This pools observations before 1970 and including and after 1970 together and 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 (coefficient 0.246, se 0.059).

How much heterogeneity is there in the effect of containerization across regions? We examine this question by estimating the specification in column (5) and drop continents one at a time. The coefficient estimates for each of these specifications are plotted in Appendix Figure A.8. The coefficient remains fairly stable and highly significant (the baseline coefficient is 0.246, while the ones that drop continents one at a time fluctuate between 0.2 and 0.3). The coefficient does become somewhat smaller when we drop North-America, which is in line with the United States being the birthplace and an early adopter of containerization. The coefficient also becomes smaller when we drop Asia, consistent with the notion that gains to containerization may have been particularly large in Asia.

The appendix contains further important robustness checks. First, we test robustness to different data construction choices. In particular, we examine different ways of treating zero shipping values, different ways of defining the depth measure for the handful of ports that are located far inland from the coastline and restricting the sample to the subset of cities that had already attained 100,000 inhabitants by 1950 to examine sample selection bias (Appendix Table A.6). The coefficient of interest remains similar in magnitude and highly significant across all these checks. Second, we examine how stable the effect of depth on shipping is depending on the buffer that we use to calculate our depth measure (Appendix Table A.7). The effect is similar across the different buffers, but becomes weaker as we move further away from the coast.[36] Having established that our depth-based measure of port-suitability predicts shipping flows after containerization, we now

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

[36] At buffers further out at sea we lose a lot of variation as many locations have a lot of depth far away from the coast.

turn to examining how this boom in shipping affected city population.

Result 2: The local causal effect of shipping on population is not distinguishable from zero. We estimate the effect of shipping on population using the following specification;

$$\ln(Pop_{it}) = \beta * \ln(Ship_{it}) + \alpha_i + \delta_t + \sum_{j=1960}^{1990} \phi_j * \ln(Pop_{i,1950}) * \mathbb{1}(Year = j) + \epsilon_{it} \quad (1)$$

where $\ln(Pop_{it})$ is the natural logarithm of population in city i at time t, and all other variables are as previously defined. The main identification challenge is that the shipping flows of a city are endogenous. Our main worry is reverse causality: fast growing cities will also witness increases in their shipping flows. Our solution is to isolate variation in shipping caused by exogenous suitability to containerization using the depth measure as an instrument for shipping. In particular, we use the binary version of our containerization treatment defined in the previous section: we interact the cross-sectional measure of depth with an indicator variable that takes the value of one in years including and after 1970. We cluster standard errors at the city level. We also report Conley standard errors for our main results in Appendix Table A.8.[37]

Table 2 contains the main regression results. In columns (1) to (6) we estimate the effects of interest using all years in the sample, while columns (7) to (10) show the estimates from the long-differenced specification. Turning first to the full panel specification, the OLS specification of equation (1) shows that the association between shipping and population is small, positive and statistically different from zero (coefficient 0.013, se. 0.005). The 2SLS estimate in column (2) shows a similarly sized coefficient but we cannot reject zero (coefficient 0.015, se. 0.049). To assess magnitudes, we report the standardized 'beta' coefficients for our effects of interest in italics underneath the estimated regression coefficients. These make clear that while the OLS may be statistically significant, the magnitudes of both the OLS and the 2SLS estimates are economically negligible. A one standard deviation increase in shipping leads to a 0.03 (OLS) or 0.035 (2SLS) standard deviation increase in population. Columns (3) and (4) show the first stage and reduced form respectively. These make clear why the results are indistinguishable from zero. While the first stage is strong (the Kleibergen-Paap F-statistic is 21.13), there is no reduced form relationship between depth and population (the reduced form coefficient is 0.004, se 0.013).[38] Columns (5) and (6) show the full time path of effects for the first stage and reduced form respectively. The time path of the fully flexible first stage was already discussed in the previous section: the effect

[37] As these are typically very close to the clustered standard errors, we only report them for the main results for easier readability of the tables.

[38] The specification here is identical to that in Table 1, but the sample size shrinks slightly as we lose those observations where population is unobserved in some years (1% of the sample).

of depth on shipping is small and indistinguishable from zero pre-containerization, and it becomes large and highly significant after the onset of containerization. The coefficients on the reduced form 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 reduced form coefficient is negative, suggesting that if anything, deeper ports were growing at a slower rate than shallower ones in some decades.

While the large standard errors typical of 2SLS estimation make a definitive answer difficult, there are several reasons why we believe that the most reasonable interpretation of our results is that shipping booms caused by containerization led to no discernible effects on population. First, as discussed above, the standardized 'beta' coefficients make clear that the magnitudes of both the OLS and the 2SLS estimates are economically negligible. Second, if we examine the long-differenced specification in columns (7) to (10), neither the OLS nor the 2SLS estimate is significantly different from zero, and both standardized beta coefficients again show an economically negligible effect. In fact, the 2SLS coefficient estimate is smaller – it is less than half the size estimated in the full panel, consistent with the fact that it was population observations from *earlier* years that drove the point estimate in the full panel specification. While the long-differenced specification has the disadvantage of using fewer observations, it has the advantage that it examines the long-run effects of the shipping boom on population, once the latter has had time to adjust.

Third, we subject the 2SLS specification to the same set of robustness checks conducted above. Appendix Table A.9 presents the results. Despite the demanding nature of these specifications, the first stage remains sufficiently strong (the Kleibergen-Paap F-statistic is always above 10) and the estimated 2SLS coefficient is never statistically different from zero. In fact, in two out of three cases, the estimated coefficient is *negative*. In particular, we estimate a negative, though statistically insignificant effect when we add the full set of coastline by year fixed effects (column 2) and when we control for initial GDP per capita by year trends (column 4). Fourth, no single continent drives this result. In Appendix Figure A.9 we plot the estimated coefficient dropping continents one at a time. The 2SLS coefficient remains close to zero and is never statistically significant. Appendix Table A.10 shows that the results are robust to various ways of treating zero shipping flows in the sample, to how we define the IV for ports further inland and to restricting the sample to cities that already reached 100,000 inhabitants in 1950.

We view the null effect on population as a surprising finding. Intuition and standard models (Coşar and Fajgelbaum, 2016; Nagy, 2018; Fajgelbaum and Redding, 2018) would both suggest that a boom in shipping should make a location more attractive for households and firms, as they can access consumers and producers more cheaply (the 'market access effect'), leading to an inflow of population. Indeed, the past literature has found that these types of positive shocks to a city's accessibility tend to lead to a boom in local population (Bleakley and Lin (2012); Brooks et al. (2019); Campante and Yanagizawa-Drott (2018).[39] While comparing the *economic* size of the effect in these papers relative to ours is difficult given the different contexts and different 'treatments', these papers all show that their effect is economically meaningful, while making the same claim with our results would be difficult.

What can explain the difference between our findings and previous work? One notable difference in our setting is the increased land intensity of of port activities induced by containerization discussed in Section 1. This may be a force that crowds out population. The large space occupied by ports, and in particular, the increased land usage necessary to adapt to containerization may have crowded out other forms of economic activity. If this force is strong enough, it could counteract the more standard, positive market access effect. In the last part of this section, we examine the extent to which we can detect the effects of this in our data. We also note that in Section 7, we will return to the question of what magnitude of an effect one would expect in our setting *in the absence* of the crowding out mechanism. We will show through the lens of our model that only by including the crowding out mechanism can we match the null population effect found in this section. Switching this mechanism off leads to a statistically and economically significant population effect – as predicted by the standard market access effect.

Result 3: Containerization increased shipping more in low rent cities. We now turn to examining whether land prices affect where port development takes place in response to containerization. We do this in the following way. To the extent that this mechanism is at work, we would expect low land-rent cities to be more attractive places for containerized ports all else equal, as the opportunity cost of port development in these locations is low. We test for this by examining the heterogeneity of the depth-shipping relationship from Result 1 using the following specification;

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

[39] The paper closest to our setting is Brooks et al. (2019), who study the effect of containerization on the population of U.S. counties located nearby. They find a positive and statistically significant effect of containerization on local population. Though the two settings are difficult to compare as we study cities around the world, we think one crucial difference is that while we examine the effects on cities, their unit of analysis is a county. As we argue below, the most likely mechanism driving the null result is that the land intensity of port technology acts as an important opposing force crowding out population. This mechanism is more likely to be detectable at the generally finer level of spatial resolution that we examine.

where $Rent_i$ is the Saiz land rent proxy for city i, and all other variables are as defined above.[40] The coefficient of interest is γ – that is, we are interested in the interaction between our depth suitability measure and the Saiz land rent proxy (interacted with the 'containerization' treatment variable that turns on in 1970). We have defined the Saiz measure such that higher values correspond to less area that can be developed, implying high land-rents. Note that this is a fully saturated specification in that we allow both depth and the Saiz measure to have their own time trend break in 1970. We plot the marginal effect of depth at different values of the Saiz measure in Figure 1 (the corresponding estimates are presented in Table 3). Consistent with the land intensive nature of containerized technology, the coefficient of interest, γ, is negative, large and statistically different from zero (coefficient -0.707, se. 0.323). Cities with exogenously deeper ports witnessed increased shipping flows after 1970, but disproportionately more so in low land rent cities.

Appendix Figure A.10 explores the heterogeneity of the result by dropping continents one at a time. The effect is consistently negative as we drop continents, though it is smaller in magnitude when we drop Asia, suggesting that the land price mechanism may have exerted a particularly strong influence in this part of the world.

We perform the same set of robustness checks for this result as for previous ones. We add coastline by year fixed effects and control for initial country GDP per capita interacted with year indicators in Appendix Table A.11. Alternatively, we treat zero shipping flows in different ways, define the depth measure for ports further inland in different ways, and restrict the sample of cities to those that already attained a population of 100,000 inhabitants in 1950 in Appendix Table A.12. The results are largely robust to these specifications, as our coefficient of interest, γ, remains negative and economically large throughout all these checks, though in two especially demanding specifications the level of significance drops below 10%.

We provide additional evidence that the land prices matter for determining the location of ports in the appendix. In Appendix Table A.13, we examine the location of ports *within* cities using information from the *World Port Index* on the geocodes of ports in our sample in 1953 and 2017.[41] We show that during this time period, ports moved further from the centroid of the city towards the outskirts, where land prices are typically lower (Duranton and Puga, 2019). This is particularly striking for the subset of cities in which a new port was built (e.g., in Sydney, Australia). In these cases, the new port was located on average 9 km further from the centroid of the city than the old port. In sum, land prices matter for determining where port development takes place.

Taking all the results of this section together, the lack of population effects following the shipping boom caused by containerization is strongly suggestive of some type of counteracting force crowding population out of the city (alongside standard forces that lead to population inflows).

[40] We report standard errors clustered at the city level, as well as Conley standard errors in curly brackets.

[41] We provide details on the data used for this exercise, and in particular, on how we calculate city centroids in Appendix C.6.

The land rent heterogeneity result suggests that the land-intensive nature of containerized technology described in Section 1 is an empirically important determinant of *where* containerized port infrastructure was developed. All else equal, port development tends to take place in cities where the land price is lower, as measured by the Saiz land rent proxy. Armed with this evidence we now turn to writing down a quantitative spatial model that captures many realistic features of port infrastructure development, including, but not limited to, the land price mechanism. In Section 5, we revisit the empirical analysis guided by the insights of the model to probe further the land price mechanism.

4 A model of cities and endogenous port development

To measure the aggregate effects of port development induced by containerization, we develop a rich and flexible quantitative general equilibrium model of trade across cities. In the model, we explicitly take into account the fact that port cities can endogenously develop their port to benefit from new port technologies. Cities facing high demand for transshipment because of their geographic position or other favorable local conditions want to invest more in developing their port. Developing the port, however, is costly as it requires scarce land that can be used for other purposes. As a result, the model captures both the benefits and costs of port development. Section 4.1 outlines the setup, while Section 4.2 discusses the qualitative predictions that the model delivers on port development and its consequences on the spatial distribution of population across cities.

4.1 Setup

The world consists of $S > 0$ cities, indexed by r or s. An exogenously given subset of cities are port cities, while the rest are non-port cities. We make the Armington assumption that each city produces one variety of a differentiated final good that we also index by r or s (Anderson, 1979). Each city belongs to one country, and each country is inhabited by an exogenous mass of workers who choose the city in which they want to live. Mobility across cities is, however, subject to frictions.

4.1.1 Workers

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

$$u_j(r) = \left[\sum_{s=1}^{S} q_j(r,s)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} a(r)\, b_j(r) \tag{3}$$

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

The dispersion of $b_j(r)$ represents the severity of cross-city mobility frictions that workers face, similar to Kennan and Walker (2011) and Monte, Redding, and Rossi-Hansberg (2018). To see why this is the case, note that if $b_j(r)$ do not vary across workers or cities, then any increase in income or amenities at r translates into a massive flow of workers towards r. On the other hand, if $b_j(r)$ are very dispersed, then workers move to the cities they prefer for idiosyncratic reasons, hence changes in economic fundamentals lead to little migration. For tractability, we assume that $b_j(r)$ is drawn from a Fréchet distribution with shape parameter $1/\eta$ and a scale parameter normalized to one. Hence, a larger value of η corresponds to more severe frictions to mobility.

4.1.2 Landlords

Each city r is also inhabited by a positive mass of immobile landlords who own the exogenously given stock of land available in the city. We normalize the stock of land available in each city to one.[42] Landlords have the same preferences over goods as workers. They do not work but finance their consumption from the revenues they collect after 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 can also use part of their land to provide transshipment services. The more land they use for transshipment services, the more the cost of transshipping a unit of a good decreases. The landlord can charge a price for the transshipment service she provides. Competition among port city landlords drives down this price to marginal cost. Hence, profits from transshipment services are zero in equilibrium.[43]

4.1.3 Production

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

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

[42] 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.

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

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 use for 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.[44] The representative firm does not internalize the effect that its employment decision has on local population. Hence, it takes $N(r)$ as given.

4.1.4 Shipping and port development

Firms in city r can ship their product to any destination $s \in S$. Shipping is, however, subject to iceberg costs: if a firm i from city r wants to ship its product over a route ρ that connects r with s, then it needs to ship $T(\rho, i)$ units of the product such that one unit arrives at s. Shipping costs consist of a component common across firms $\bar{T}(\rho)$, as well as a firm-specific idiosyncratic component $\epsilon(\rho, i)$ that is distributed iid across firms and shipping routes:[45]

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

For tractability, we assume that $\epsilon(\rho, i)$ 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,

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

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

Certain shipping routes involve land shipping only (*land-only*), while others involve a combi-

[44] 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).

[45] The assumption of idiosyncratic shipping cost shifters follows Allen and Atkin (2016) 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 (2020), finding optimal shipping flows is computationally more demanding.

nation 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, $d(r, s)$:

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

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

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

where $\phi_{\varsigma}(d(r, p_0))$ corresponds to the overland shipping cost between the origin and the first port en route p_0, and $\phi_{\varsigma}(d(p_M, s))$ corresponds to the overland shipping cost between the last port en route p_M and the destination. $\phi_{\tau}(d(p_m, p_{m+1}))$ denotes the sea shipping cost between ports p_m and p_{m+1}, a function of the minimum sea distance between the two ports, $d(p_m, p_{m+1})$. Finally, $O(p_m)$ denotes the price that the firm needs to pay for transshipment services in port city p_m.[46]

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

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

where $\nu(p_m)$ is an exogenous cost shifter capturing the fundamental efficiency of port p_m, $\psi(F(p_m))$ is a non-negative, strictly decreasing and strictly convex function of $F(p_m)$, the share of land allocated to the port, and $Shipping(p_m)^{\lambda}$ captures congestion externalities arising from the fact that handling one unit of cargo becomes more costly as the total amount of shipping, $Shipping(p_m)$, increases for a given port size.[47] As each port city landlord is atomistic, she takes the price of transshipment services $O(p_m)$ and the total port-level shipping $Shipping(p_m)$ as given when choosing $F(p_m)$. Moreover, perfect competition among port city landlords ensures that the price of transshipment services is driven down to marginal cost and therefore

$$O(p_m) = [\nu(p_m) + \psi(F(p_m))] \, Shipping(p_m)^{\lambda} \tag{4}$$

[46]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.

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

in equilibrium.

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

4.1.5 Equilibrium

In equilibrium, workers choose their consumption of goods and residence to maximize their utility, taking prices and wages as given. Landlords choose their consumption and land use to maximize their utility, taking prices, land rents and shipping flows as given. Firms choose their production of goods, employment and land use to maximize their profits, taking prices, land rents and wages as 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. Appendices B.1 and B.2 provide a formal definition and characterization of the equilibrium.

4.2 Predictions of the model

In equilibrium, the share of land allocated to the port in port city r is the solution to the equation

$$-\psi'\left(F\left(r\right)\right)=\frac{R\left(r\right)}{Shipping\left(r\right)^{1+\lambda}} \tag{5}$$

where $R\left(r\right)$ denotes land rents in city r, given by

$$R\left(r\right)=\frac{1-\gamma}{\gamma}\frac{w\left(r\right)N\left(r\right)}{1-F\left(r\right)} \tag{6}$$

such that $w\left(r\right)$ is the wage in city r.[48] As the left-hand side of equation (5) is decreasing in $F\left(r\right)$ by the convexity of ψ, we have the following two propositions.

Proposition 1. *Land allocated to the port is increasing in the amount of shipping flows.*

Proposition 1 is the consequence of two forces in the model. The first is economies of scale in transshipment technology: as shipping flows increase, it becomes profitable to lower unit costs by allocating more land to the port. The second force is congestion: an increase in shipping flows makes landlords allocate more land to the port to palliate congestion.

Proposition 2. *Land allocated to the port is decreasing in land rents.*

[48] The derivation of equations (5) and (6) is included in Appendix B.2.

Proposition 2 highlights that the cost of adopting containerized technologies differs across cities. Cities that have high land rents do not allocate much land to the port as the opportunity cost of land is very high. As a result, everything else fixed, port development primarily takes place in low-rent cities, consistent with what we document in the data.

Finally, the model delivers the spatial distribution of population $N(r)$ as the solution to the following equation:

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

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

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

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.[49]

How is the population of a port city affected by the development of its port? Our last proposition shows that the net effect on population is the outcome of two opposing forces: the *market access effect* that increases the population of the city, and the *crowding-out effect* that leads to a decrease in the city's population.

Proposition 3. *An increase in the share of land allocated to the port in city r, $F(r)$, decreases shipping costs $\mathbf{E}[T(r,s)]$, thus increasing $MA(r)$. Everything else fixed, an increase in $MA(r)$ increases the population of the city (market access effect). At the same time, holding $MA(r)$ fixed, an increase in $F(r)$ decreases the share of land that can be used for production, $1-F(r)$, thus decreasing the population of the city (crowding-out effect).*

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

Proposition 3 sheds light on the fact that, to measure the net effect of port development, it is essential to consider both its benefits and its costs. On the one hand, port development lowers shipping costs. On the other hand, it requires scarce local land that needs to be reallocated from other productive uses. The model, and equation (7) in particular, provide a structure that allows

[49] Appendix B.2 provides the derivation of equation (7).

us to capture these opposing forces. The next section is aimed at looking for evidence on these opposing forces in the data.

5 Empirical evidence for the model's mechanisms

In this section, we examine whether there is empirical evidence for the two opposing model forces through which port development affects local city population. Equation (7) shows that port development has a positive effect on population through lowering transshipment costs which will increase the market access term. However, holding market access fixed, port development has a *negative* effect on population as it crowds out non-port activities.

Due to the lack of time-varying data on port sizes, we cannot directly take equation (7) to the data. However, we can estimate a simplified version of the equation to understand whether we can disentangle the positive and negative effects of port development predicted by the model. Based on this reasoning, we estimate the following relationship, referred to as *model-inspired empirical specification*:

$$\ln(Pop_{it}) = \phi_1 * \ln(Ship)_{it} + \phi_2 * \ln(MA)_{it} + \alpha_i + \delta_t + \epsilon_{it} \tag{8}$$

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

Constructing the market access term requires us to estimate time-varying bilateral trade costs $T_t\,(i,s)$ between origin and destination. As in the model, we assume that these bilateral costs consist of a combination of three possible components: first, the cost of shipping overland; second, the cost of sea shipping; and third, the cost of transshipment at seaports. Armed with these, we use the fast marching algorithm to calculate the lowest overall shipping cost between any given pair of cities. Following Allen and Arkolakis (2014), we assume that overland shipping costs ϕ_ς and sea shipping costs ϕ_τ take the form

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

where d is (point-to-point) distance traveled. We take the values of t_ς and t_τ from the road and sea shipping cost elasticities estimated by Allen and Arkolakis (2014).[50]

There are no readily available measures of transshipment costs that we are aware of. To con-

[50] Allen and Arkolakis (2014) also allow for costs of inland and sea shipping that are fixed with respect to distance. However, they set the fixed costs of road shipping to zero. In the case of sea shipping, our aim is to define transshipment costs incurred at the seaport in a broad sense, such that they include any cost that is not a function of shipping distance, such as the fixed costs of sea transportation.

struct these, we use the following approach. Both the model and the transportation literature on ports argue that there are important economies of scale in port technologies rendering larger ports more cost-efficient (Rodrigue, 2016). We use estimates of port costs, available for a subset of our ports from Blonigen and Wilson (2008), to estimate the empirical relationship between port costs and shipping flows at the port level in our data using a simple linear OLS specification.[51] Consistent with economies of scale in shipping, we find a negative and statistically significant association between port costs and the size of shipping flows. We use the estimated coefficient from this regression to predict port efficiency for all the ports in our data for each decade.[52] Note that changing transshipment costs are the only source of time series variation in our estimated trade costs.

The model-based measure of market access requires taking a stand on the values of the parameters η, σ, γ and α. Table 5 contains the parameter values we use and their source. As we use the same values when taking the full model to the data, Section 6.2 discusses the calibration of all structural parameters in detail.

Both regressors in the model-inspired specification (8) are potentially endogenous. Identifying ϕ_1 and ϕ_2 thus requires two sources of exogenous variation. We use the baseline measure of depth in the vicinity of the port as an instrument for shipping, as explained in Section 3. We use an exogenous population-growth shifter based on regional climate to construct an instrument for market access. This IV is based on insights from the urban economics literature which has found that people have moved to places with warm winters over the course of the 20th century – a phenomenon attributed to the invention of air conditioning (e.g., Oi (1996); Rappaport (2007)).

In order to implement this in our setting, we use the average number of frost free days, $frostfree_i$, during the years between 1961-1990 in each city to predict population growth during our time period.[53] Notice that this second instrument uses a source of exogenous variation that is orthogonal to port depth, as the number of frost free days and port depth (both residualized and un-residualized) are uncorrelated with each other.[54] In order to predict population, we estimate the following specification:

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

where β_k estimates the effect of warmer winters on population in each decade, α_i denotes city-

[51] See Appendix C.10 for more details.

[52] Details on this estimation are provided in Appendix Tables A.14 and A.15. Note that we predict the port cost for the full set of 2,145 ports in our data set. These include the set of ports for which we do not have population data. This is in order to allow for the most realistic trade routes. See Appendix Table A.2 which shows the breakdown of different types of ports in the sample.

[53] Appendix C.11 describes the data on the number of frost free days.

[54] The correlation between the number of frost free days and unresidualized depth is 0.04 (p-val: 0.40), and the correlation between the number of frost free days and residualized depth is -0.02 (p-val: 0.68).

specific fixed effects, and δ_{ct} allows for the full set of country by year fixed effects. Inclusion of these implies that we only use *within-country* variation in climatic conditions when estimating the effect of frost-free days on population growth. We do this to address the concern that climatic conditions vary across regions in ways that may correlate with unobserved drivers of population growth, confounding our estimates of interest. Appendix Table A.16 shows the result of this estimation and presents some robustness checks. To construct our second instrument for market access, we predict population for each city – year pair based on the estimated effects of frost free days and the estimated city fixed effect (we do not use the estimated country-year fixed effects to predict population):

$$\widehat{\ln(Pop)}_{it} = \sum_{k=1960}^{1990} \widehat{\beta}_k * frostfree_i * \mathbb{1}(Year = k) + \widehat{\alpha}_i$$

Using these predictions for city-level population, we define our second instrument as follows:

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

where $T_{1950}(i,s)$ is the transport cost between cities i and s in 1950. We hold bilateral transport costs fixed throughout all years in order to make sure that potentially endogenous changes in trade costs over time are not used in the instrument.

A final question is what sample should be used for the estimation of equation (8). The model applies to all cities, regardless of whether they are port or non-port cities. However, as port development opportunities are only available for port cities, the crowding-out mechanism will only be relevant for these cities. Moreover, the IV used to identify ϕ_1 (port depth, as described in section 3) is only defined for port cities. Accounting for non-port cities, however, is important as the general equilibrium effects of port development elsewhere will impact population and trade costs to these cities, and hence affect port cities. For this reason, while the specifications are estimated on the set of *port* cities in our dataset, the market access of port cities is calculated using the full set of (port and non-port) cities.

Table 4 presents the estimation results. Columns (1) and (2) report the baseline reduced form OLS and 2SLS estimates for comparison. Columns (3) and (4) add the measure of market access as a control. The OLS estimate in column (3) shows a very small negative effect of shipping on population relative to column (1) that is not distinguishable from zero. Column (4) shows the 2SLS specification. Consistent with the predictions of the model, once we control for market access, shipping has a negative, statistically significant effect on population.[55] The instruments yield a

[55] As expected, market access has a significant positive effect on population. It is difficult to compare the size of the market access effect to existing estimates (Donaldson and Hornbeck, 2016; Jedwab and Storeygard, 2020; Maurer and Rauch, 2020) because different papers construct market access in different ways. Jedwab and Storeygard (2020) are

combined Kleibergen-Paap F-statistic of 9.63 which is just below the often recommended value of 10; however, it is larger than the critical value of 7.03 that the Stock-Yogo weak ID test suggests for 10% maximum bias (Stock and Yogo, 2002).[56] Columns (5) and (6) report the first stages of the regression. Reassuringly, depth is a strong predictor of shipping, while the market access IV predicts market access strongly. Appendix Table A.17 shows that the pre-trends check with respect to depth holds (for both first stages) in this more complex specification that adds market access.[57]

We test the robustness of this result in a number of ways. Appendix Table A.18 shows that the results are remarkably robust to dropping cities in the close vicinity of the city in the market access IV, suggesting that much of the identifying variation is coming from population movements further away from the city itself. Appendix Table A.19 shows that the sign of the effects are robust to the same set of controls used in Section 3, though in the case of these demanding specifications, we don't always retain statistical significance at 10%.

In summary, these results show strong support for the model mechanisms. On the one hand, new port technologies improve a location's market access, drawing population in. This market access effect is well known in the literature and has been found to be present in different contexts (Donaldson and Hornbeck, 2016; Jedwab and Storeygard, 2020; Maurer and Rauch, 2020). On the other hand, there is a marked, negative, direct effect of shipping on economic activity, consistent with the crowding-out effect of port development. We conclude that this lends well-identified evidence for the model mechanism. In the next section, we therefore turn to taking the full model to the data.

6 Taking the model to the data

We take the full structure of the model to the data in this section. This allows us to estimate the aggregate effects of changing port technologies in Section 7.

Taking the model to the data consists of three steps. In the first step, we calculate inland and sea shipping costs across cities and choose a functional form for endogenous transshipment costs as a function of land use, $\psi(F)$. In the second step, we choose the values of the model's seven structural parameters. In the last step, we back out the values of unobserved city fundamentals (amenities, productivities and exogenous transshipment costs) using a cross section of observed city characteristics: population, shipping flows and GDP. Below, we describe each of these steps in detail.

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

[56]With the usual caveat that Stock and Yogo (2002) values have been derived only for i.i.d. errors, whereas we allow for autocorrelated or spatially correlated standard errors.

[57]As there is no similar 'pre-treatment period' for the market access IV, it is not possible to conduct a similar exercise for this IV.

6.1 Calculating shipping costs

We follow our strategy outlined in Section 5 to calculate inland and sea shipping costs across cities[58] as a function of distance d, assuming

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

and setting the elasticities t_{ς} and t_{τ} to the corresponding estimates in Allen and Arkolakis (2014).[59]

We also need to choose endogenous transshipment costs as a function of the share of land allocated to the port (*port share*, F), $\psi(F)$. The existing literature provides us with little guidance on this, as ours is the first paper that argues for the relevance of this relationship in a quantitative trade and geography framework. Hence, our goal is to keep the functional form of ψ as simple as possible. That said, the functional form needs to satisfy our theoretical restrictions ($\psi \geq 0, \psi' < 0, \psi'' > 0$) and needs to be numerically tractable in the model inversion and counterfactual simulations. In particular, the range of ψ' should ideally span the entire $(-\infty, 0)$ interval over its domain $(0, 1)$, as otherwise it would be potentially impossible to obtain port shares that rationalize the GDP and shipping data in every port city from equations (5) and (6). One simple function that satisfies all these restrictions is

$$\psi'(F) = 1 - F^{-\beta} \tag{9}$$

where we restrict $\beta > 0$ to guarantee $\psi' < 0$. We can then obtain ψ by integrating equation (9) as

$$\psi(F) = \frac{F^{\beta} + (\beta - 1)^{-1}}{F^{\beta - 1}} + \kappa$$

where we restrict $\kappa \geq \bar{\kappa} = -\left[1 + (\beta - 1)^{-1}\right]$ to guarantee $\psi \geq 0$.[60]

6.2 Choosing the values of structural parameters

We also need to choose the values of the model's seven structural parameters. On the production side, we take the estimate of the strength of agglomeration externalities, $\alpha = 0.06$, from Ciccone and Hall (1993). This estimate has performed well in the literature for various countries and time periods. $\alpha = 0.06$ implies that doubling city size increases city productivity by 6%. Still on the production side, the expenditure shares on labor and land equal γ and $1 - \gamma$, respectively. Unfortunately, we are not aware of any study that measures the land share for the entire world.

[58] We have 553 port and 2,083 non-port cities in our data. For details, see Section 2.

[59] See Section 5 for details on these estimates.

[60] As total transshipment costs in city r equal $[\nu(r) + \psi(F(r))]\, Shipping\,(r)^{\lambda}$, κ is isomorphic to a uniform shifter in exogenous port costs $\nu(r)$ and therefore cannot be identified separately from them. Thus, we set κ to its theoretical lower bound $\bar{\kappa}$ without loss of generality.

Thus, we base our benchmark value of γ on Desmet and Rappaport (2017), who estimate a value of 0.10 for the difference between the land share and the agglomeration elasticity in the United States between 1960 and 2000, a period that corresponds to our sample period. Given we set $\alpha = 0.06$, this suggests choosing $\gamma = 0.84$.[61]

On the consumption side, we have two structural parameters: the migration elasticity, which we set to $\eta = 0.15$ based on Kennan and Walker (2011), and the elasticity of substitution across tradable final goods, which we set to $\sigma = 4$ based on Bernard, Eaton, Jensen, and Kortum (2003).

Finally, there are three structural parameters that influence shipping costs. 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 $\lambda = 0.074$. Table 5 summarizes the calibration of our structural parameters.

The last structural parameter to choose is β from the endogenous transshipment function. Given the role that this parameter plays in driving the relationship between the value of shipping flows and the port share through equation (5), we calibrate it to match the correlation between these two variables in the data.[62] We use the port share data constructed for seven cities that was described in Section 2. The correlation between shipping and port share for these seven cities is 0.474.

In the model, we compute the correlation between shipping and port share in the following way. First, for each port city, we numerically solve equations (5) and (6) for the port share that rationalizes shipping flows, $Shipping\,(r)$, and city GDP, $\gamma^{-1} w\,(r)\,N\,(r)$. As we explain in Appendix B.3, our theoretical restrictions on ψ' guarantee that this procedure identifies a unique port share $F\,(r) \in (0,1)$ for each port city. Next, we calculate the correlation between $Shipping\,(r)$ and $F\,(r)$ for our set of port cities.

Under higher values of β, the endogenous port development mechanism plays a stronger role in the model. This is because, under higher β, the endogenous transshipment cost function is more responsive to changes in the port share:

$$\frac{d\,|\psi'\,(F)|}{d\beta} = -F^{-\beta}\log\,(F) > 0$$

Hence, everything else fixed, landlords have an incentive to increase the port share further if β is high. As a consequence, we expect a stronger correlation between shipping and port share under

[61] Another advantage of using the land share estimate by Desmet and Rappaport (2017) 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.

[62] To calculate this correlation, we first transform the number of ships, which is what we directly observe in the data, into the value of shipments, which is what enters equation (5). This procedure is described in detail in Section 6.3.

higher values of β. This is precisely what we find. Appendix Figure A.11 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.[63] Hence, we use this value of β in our baseline calibration.[64]

6.3 Recovering post-containerization fundamentals

In the final step of taking the model to the data, 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 city amenities $a(r)$, 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. Since this year is after the advent of containerization, the counterfactual we will simulate in Section 7 to estimate the aggregate effects of containerization will *roll back*, or undo, the containerization shock. Hence, the 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 $V = 364$.[65]

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, as explained in

[63]Instead of calculating the model-implied correlation over the entire set of port cities, we can compute it for the same set of seven port cities where we observe the port share. Reassuringly, for $\beta = 0.031$, this gives us a correlation of 0.463, essentially identical to the one found for the whole set of port cities.

[64]In Section 7, we investigate robustness of the aggregate gains from containerization to alternative values of β.

[65]As not all our port cities have a 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.

Section 6.2.

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 B.3.[66]

7 The aggregate effects of containerization

To measure the aggregate effects of containerization, we use our model to conduct a counterfactual in this section. As we took the model to post-containerization (1990) data in Section 6, our counterfactual involves *rolling back containerization*: i.e., changing port technologies back to pre-containerization technologies. At the same time, we keep all other fundamentals of the model (city amenities and productivities, inland and sea shipping costs and country populations) fixed. Hence, comparing the 1990 equilibrium to the counterfactual equilibrium allows us to measure the aggregate effects of containerization on the world economy.

7.1 Counterfactual: Rolling back containerization

How do we roll back containerization? As we argued in Section 1, containerization had two major effects on port technologies. First, it decreased transshipment costs, especially in deep ports due to increased ship sizes. Second, it increased the land intensity of transshipment. In our counterfactual, we incorporate these aspects of containerization by changing transshipment technology in the following way. To capture the higher transshipment costs of pre-containerization technologies, we increase exogenous transshipment costs $\nu(r)$ uniformly across ports relative to the 1990 values of these costs. To capture the fact that containerization made port depth relevant for transshipment, we offset the negative relationship between $\nu(r)$ and depth that we observe in 1990. Finally, to capture the lower land intensity of pre-containerization technologies, we decrease the shape parameter of our endogenous transshipment cost function, β.[67]

As discussed in Section 6.2, a decrease in β makes the endogenous transshipment cost function less responsive to changes in the port share, $F(r)$. Hence, under lower values of β, port city landlords have less incentive to increase $F(r)$. As a result, port sizes will be generally smaller in the counterfactual. To choose the value of the parameter in the counterfactual, β_{CF}, we use the well-documented evidence on New Orleans described in Section 1. In particular, we argue

[66] 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.

[67] Appendix B.4 describes the details of how we solve for the equilibrium of the model under these new fundamentals.

in Section 1 that the size of the port of New Orleans increased by 75% due to the technological aspects of containerization. In our model, this means that the port share of New Orleans would have increased by 75% if we *kept the non-technological determinants of the port share*, i.e., shipping and land rents, fixed:

$$\frac{F\left(\text{New Orleans}\right)}{\hat{F}\left(\text{New Orleans}\right)} - 1 = 0.75 \tag{10}$$

where $F\left(\text{New Orleans}\right)$ is the port share of New Orleans in 1990, given by

$$-\left[1 - F\left(\text{New Orleans}\right)^{-\beta}\right] = \frac{R\left(\text{New Orleans}\right)}{Shipping\left(\text{New Orleans}\right)^{1+\lambda}} \tag{11}$$

which we obtain by combining equations (5) and (9), and $\hat{F}\left(\text{New Orleans}\right)$ is the port share implied by the *same* rents and shipping but shape parameter β_{CF}:

$$-\left[1 - \hat{F}\left(\text{New Orleans}\right)^{-\beta_{CF}}\right] = \frac{R\left(\text{New Orleans}\right)}{Shipping\left(\text{New Orleans}\right)^{1+\lambda}} \tag{12}$$

To back out β_{CF}, we first solve equation (11) for $F\left(\text{New Orleans}\right)$. Next, we solve equation (10) for $\hat{F}\left(\text{New Orleans}\right)$. Finally, we solve equation (12) for β_{CF}. This procedure yields $\beta_{CF} = 0.021$.

To offset the relationship between exogenous transshipment costs and depth, we first run the regression

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

on our sample of port cities, where $\nu\left(r\right)$ is the exogenous transshipment cost of city r recovered in Section 6.3, and $Depth\left(r\right)$ is our residualized depth measure, defined in Section 3. In line with the fact that depth lowers transshipment costs after containerization, we find $\widehat{\omega_1} = 0.048$ (se. 0.025, p-value 0.053). We then undo the dependence of exogenous transshipment costs on depth by adding $\widehat{\omega_1} * Depth\left(r\right)$ to $\log \nu\left(r\right)$.[68]

Finally, we incorporate the overall reduction in transshipment costs due to containerization by increasing $\log \nu\left(r\right)$ uniformly across ports. In particular, we add a constant $\nu_{CF} > 0$ to all $\log \nu\left(r\right)$ such that total transshipment costs are 25% higher in our counterfactual than in 1990. We choose this number in the following way. Rodrigue (2016) estimates that containerization led to an overall 70% to 85% reduction in maritime transport costs by 2010.[69] What fraction of this cost reduction

[68] To avoid outliers influencing the results of this step, we trim the values of $\nu\left(r\right)$ at 0.01 before estimating ω_1. The inversion algorithm assigns very small ν's to some cities. Due to lack of machine precision in the inversion algorithm for very small values of ν's, very small differences in ν's may be exaggerated greatly when taking logs of ν's in the estimation. This affects 19 port cities, for which we identify $\nu\left(r\right)$ below 0.01. As these 19 port cities are deeper than average, we obtain a slightly higher regression coefficient, $\widehat{\omega_1} = 0.056$, without the trimming.

[69] Rodrigue (2016, p. 117) states: "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

happened prior to 1990? We know that container ship sizes increased prior to 1990 by 36% of the overall increase until 2010.[70] Assuming that cost reductions were proportional to ship size increases, approximately 36% of the 70% cost reduction must have happened before 1990 (using the more conservative end of Rodrigue's estimate). This gives us a 25% decrease in 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 $\nu_{CF} = 0.280$.

Overall, according to our simulation, these changes in transshipment technology lead to an increase in the international trade to world GDP ratio by 4.8 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-third* of the overall increase in trade to world GDP during these three decades.

The land occupied by ports (i.e., the port share) increases in most port cities from the counterfactual to the 1990 equilibrium. Port shares become larger for two reasons. First, the increase in β increases the incentive to invest more land in port development, mimicking the changing land-intensity of port technologies caused by containerization. Second, the reduction in trade costs leads to increased demand for shipping, encouraging yet more investment in port development. Appendix Figure A.12 presents the full distribution of port share changes across cities. The median change is 2 percentage points, while the 5th percentile is -0.1 pp and the 95th percentile is 32 pp. As a comparison, in our data for New Orleans, we find that the port share increased by 0.29 pp between 1960 and 1990, which puts this city at the 40th percentile of our model-implied port share change distribution. As New Orleans was *not* among the more prominent adopters of containerized technology[71], its position in the model-implied distribution lends additional credibility to the port share changes implied by the model simulation.

7.2 Test of the model: The reduced-form effects of containerization

In this section, we show that our quantified model can replicate the two key reduced-form facts related to containerization estimated in Section 3: the fact that the local causal effect of shipping on population is indistinguishable from zero, as well as the fact that containerization increased shipping more in low-rent cities. As we do not feed these reduced-form results into either the model calibration or the counterfactual, we view these as tests of the model.

To examine the local population effects of shipping in the model, we consider the specification

$$\Delta \ln(N_i) = \beta * \Delta \ln(Shipping_i) + \Delta\epsilon_i \tag{13}$$

reduction from 10% to 1.5% equals an 85% cost reduction.

[70]This can be seen in Appendix Figure A.2, which is based on data from Haworth (2020).

[71]For example, New Orleans built its first dedicated containerized facility relatively late, in 1975.

where Δ denotes the change in a variable from the counterfactual to the 1990 equilibrium. We instrument the change in shipping with residualized port depth. Thus, specification (13) is analogous to the long-differenced version of specification (1). The only difference is that while we ran the long-differenced version of equation (1) on 1950 and 1990 data, we run equation (13) on the model-simulated counterfactual and 1990 data.

Table 6 presents the results of this exercise. Column (1) replicates the estimates obtained from the long-differenced specification in the data (which is, therefore, identical to column (8) of Table 2). Column (2) presents the estimated coefficient from the model-simulated data. Recall that shipping had no discernible causal effect on local city-population in the data. We find the same null result in the model-simulated data. As the size of the shock at the city level is potentially different between the model and the data, the magnitudes of the estimated coefficients are not directly comparable between columns (1) and (2). However, we report the corresponding standardized coefficients, which are comparable, in italics. These demonstrate that a one standard deviation increase in shipping translates into a negligible (0.006 standard deviation) population gain in the model as well as in the data (0.022) – neither estimate is statistically or economically significant.[72]

In Section 3, we interpreted the estimated coefficient as a surprising result. A boom in shipping due to containerization does not translate into population gains, suggesting that there is a force crowding out population alongside the standard positive market access effect. Our model has such a force: the crowding-out effect of increased land use caused by port development. Column (2) confirms that this force is sufficient to eliminate the positive local population effect of increased shipping due to the market access effect, making the causal effect of shipping economically and statistically insignificant. In other words, we can obtain a crowding-out effect in the model that is strong enough to replicate the zero population effects of shipping observed in the data. This is true despite the fact that, as we mentioned in Section 7.1, the model-implied increases in land use are not particularly large.

Armed with this evidence, we now examine whether containerization makes shipping activity reallocate toward low-rent cities in the model. To this end, we consider the specification

$$\begin{aligned} \Delta \ln(Shipping_{it}) &= \beta * Depth_i + \gamma * Depth_i * \ln(R_{i,CF}) \\ &\quad + \eta * \ln(R_{i,CF}) + \Delta\epsilon_{it} \end{aligned} \tag{14}$$

where t is one of the two time periods: the 'counterfactual period' or 1990. This specification is the long-differenced version of specification (2), which we ran on the full panel. $Depth_i$ is our depth measure residualized on population in 1950, as in the data. The difference is that, while we had to rely on a proxy of city-level rents in specification (2), we can use model-implied (pre-

[72] We discuss columns (3) and (4) of Table 6 in section 7.3.

containerization) rents $R_{i,CF}$ in specification (14). Our coefficient of interest is γ, that is, the interaction between port depth and rents.

We evaluate the coefficient of interest, γ, at different values of log rents $\ln(R_{i,CF})$ in Figure 2.[73] As the figure shows, the effect of land rents on shipping is negative, large and statistically significant, as is the case in the data. Thus, the model can successfully replicate the finding from Section 3 that containerization increased shipping more in initially low land rent cities. This provides further evidence that the land price mechanism is present in the model not only in a qualitative sense (as we showed in Section 4.2), but it is a significant driver of where port development takes place. In summary, we view these results as providing validation for the model's ability to capture the main forces affecting port development. In the next section, we therefore turn to discussing our estimates of the aggregate effects.

7.3 The aggregate welfare effects of containerization

We estimate that aggregate world welfare increased by 3.95% as a result of containerization.[74] The welfare gains from containerization stem from a combination of three factors in the model: lower shipping costs, which increase welfare; the increased cost of land use, i.e., the *resource costs* of containerization, which lower the gains; and the gains from increased specialization of cities in port or non-port activities, i.e., the *specialization gains* from containerization.

To assess the quantitative importance of each of these margins, we develop two simple benchmark models that will allow us to isolate the three mechanisms at work. 'Benchmark 1' is closest to a standard model, as it assumes that transshipment costs are *exogenous* and *free* – that is, land is solely used for the production of the city-specific good. Thus, the welfare gains from containerization only stem from shipping cost reductions in this benchmark model. 'Benchmark 2,' on the other hand, allows for both exogenous and endogenous transshipment costs, such that endogenous transshipment costs depend on land use, as in our baseline model. However, we restrict land use to be identical across port cities (and equal to the average port share in our baseline).

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.

To implement the decomposition of the aggregate welfare effects, we follow a procedure sim-

[73] Appendix Table A.20 shows the corresponding estimates.

[74] We define the change in aggregate world welfare as the average of changes in country-level welfare between the counterfactual and the 1990 equilibrium, weighted by country population. Within each country, labor mobility equalizes welfare across cities, as in Redding (2016). However, we do not allow for mobility across countries, hence different countries experience different welfare effects. We discuss these country-level effects in Section 7.4.

ilar to the one described in Section 6.3 to take Benchmark 1 and Benchmark 2 to our 1990 data. Next, we conduct the containerization counterfactual in each benchmark model. In particular, we conduct the counterfactual such that the world trade to GDP ratio changes to the same extent ($+4.8\%$) in each benchmark as in our baseline model. Hence, differences in the welfare effects across the models do not stem from trade changing to a different extent in one versus the other.[75]

We find that containerization leads to welfare gains of 4.15% in Benchmark 1. In other words, the gains from the shipping cost reduction caused by containerization amount to 4.15% of world welfare. In Benchmark 2, the gains from containerization reduce to 3.60%. The difference between Benchmark 1 and Benchmark 2, 0.55 percentage points, captures the resource costs of containerization. These costs are sizeable: they eat up as much as 13.3% of the gains from the shipping cost reduction. Finally, the difference between Benchmark 2 and our baseline model, 0.34 percentage points, captures the specialization gains from containerization. Note that these gains are able to offset about 63% of the resource costs of containerization, but they do not fully compensate for all the costs. Based on this exercise, relative to a standard model in which transport cost reductions are exogenous and free, both model mechanisms – the resource cost and the endogenous specialization effect – lead to quantitatively meaningful effects on welfare. As the resource cost effect is larger than the endogenous specialization effect, on net the gains in our model end up being somewhat *smaller* than what a standard model would predict.

Besides using the two benchmarks for the decomposition of aggregate welfare effects, we can also use them to provide a further test of whether it is indeed our endogenous crowding-out mechanism that leads to the null effect of shipping on population in the model. To this end, we estimate the causal effect of shipping on population – equation (13) – in the two benchmarks, and contrast them with the baseline model. Columns (3) and (4) of Table 6 report the results for Benchmark 1 and Benchmark 2, respectively. Unlike in our baseline model, shipping leads to a significant increase in city population in both benchmarks. This is intuitive: while better market access draws people into the city in all three models, increased land use in transshipment does not have a differential impact on city population in the benchmarks.[76] To compare the magnitudes of the estimated coefficients, we report the standardized coefficients in italics. These demonstrate that a one standard deviation increase in shipping translates into substantially larger (0.124 and 0.14 standard deviation) increase in population in the benchmarks than in the baseline model (0.006) or in the data (0.022). This underscores that the crowding out effect is driving the zero local population effect of shipping in the model. It also points to the fact that the crowding-out effect is

[75] We provide a detailed description of each benchmark model, the procedure of taking them to the data and the procedure of conducting the counterfactual in them in Appendix B.5.

[76] In Benchmark 2, land used for transshipment increases equally across port cities. Hence, land used for transshipment does not react endogenously to shipping, leading to no differential impact on the population of cities with different changes in shipping. In Benchmark 1, no land is used for transshipment by assumption.

sizeable – not just in terms of the effect it has at the aggregate level, but also in terms of its local effect.

In Table 7, we examine the sensitivity of our headline aggregate welfare result to different values of the containerization shock and some alternative modeling choices. In rows (2) and (3), we use higher and lower values of our transshipment cost parameter β, respectively. In rows (4) and (5), we use alternative values of our counterfactual β: one that implies a smaller (65%) increase in the port share of New Orleans, and one that implies a larger (85%) increase. As expected, a smaller increase in land use leads to slightly higher welfare gains from containerization. In row (6), we do not offset the relationship between exogenous transshipment costs and port depth in the counterfactual. In rows (7) and (8), we choose ν_{CF} to target different (30% and 20%, respectively) changes in total transshipment costs. Finally, in row (9), we take the model with monopolistic competition, presented in the Appendix B.7, to the data. The key difference relative to our baseline setup is that port activity involves positive profits in the monopolistic competition model. The welfare gains from containerization are fairly stable across these different specifications with welfare effects ranging from 3.3 – 4.6. We conclude that the estimated effects would be similar had we chosen slightly different parameter values for the shock.

7.4 Country-level effects

In this section, we investigate the country-level welfare effects of containerization implied by the model. Appendix Figure A.13 plots the distribution of the welfare gains from containerization across countries. As the figure demonstrates, these gains vary substantially around the worldwide average (3.95%). For instance, 26% of countries see welfare gains below 2%, while 29% experience gains above 10%. Our goal in this section is to understand the sources of these cross-country differences in the gains from containerization.

Plotting the welfare gains from containerization against country GDP per capita prior to containerization, one can observe a negative relationship between these two variables (Appendix Figure A.14).[77] That is, initially poorer countries benefited more from containerization on average. Column (1) of Table 8 contains the estimated (standardized) coefficient in a regression of the welfare gains on country GDP per capita. According to this estimate, a one standard deviation lower GDP per capita is associated with 0.161 standard deviation higher gains. That is, the effect is significant not only statistically, but also economically.

What allows poorer countries to benefit more from containerization? In Table 8, we regress the country-level welfare gains from containerization on a set of country-level covariates. From column (1) to column (5), we gradually add covariates until the negative effect of GDP per capita is

[77] We calculate country GDP per capita as the average of counterfactual GDP per capita across cities, weighted by counterfactual city population. The results are similar if we use actual country GDP per capita prior to containerization (1960).

fully soaked up by them. Column (5) shows that this is achieved by including the share of port cities in the total set of cities, country population, average productivity and average pre-containerization market access on the right-hand side.[78] Finally, in column (6), we drop GDP per capita but keep the other covariates. This is our preferred specification as GDP per capita and average productivity are highly correlated (correlation coefficient 0.95), causing potential multicollinearity problems if both variables are added simultaneously. Market access is also positively correlated with GDP per capita, though much less strongly than productivity (correlation coefficient 0.05). Overall, these findings suggest that poorer countries gained more from containerization because they had lower average productivity (in non-port activities) and worse pre-containerization market access.[79]

The result that countries with lower average productivity and worse pre-containerization market access benefit more from containerization is intuitive. First, note that productivity refers to productivity in non-port activities, suggesting that lower-productivity countries have a comparative disadvantage in these sectors, and a comparative advantage in port-activities. Thus, it is no surprise that they benefit more from a positive shock to the shipping sector. Second, note that the containerization shock lowers shipping costs, which should benefit more peripheral countries (those with initially worse market access). As shipping costs decrease, peripheral countries improve their market access dramatically, while central countries already had good market access to begin with. Moreover, we expect this force to be amplified by endogenous port development. If peripheral countries are poorer and have lower land rents, then they should also benefit from the possibility of developing their ports at a low cost, which leads to an additional reduction in their shipping costs.

To examine whether this intuition for the role of market access is correct, we compare the relationship between market access and the welfare gains from containerization in our two benchmark models. Recall that these benchmarks do not feature endogenous port development at the city level. Thus, our reasoning suggests that they should feature a negative but weaker relationship between market access and the gains from containerization than our baseline model. This is precisely what we find. Appendix Figure A.15 plots the gains from containerization against pre-containerization market access in our baseline model, in Benchmark 1, and in Benchmark 2. Though each model features a negative relationship between these two variables, the relationship is the strongest in our baseline model.[80]

[78]To isolate the part of market access that is a function of geography alone, the market access variable used for this exercise does not include population. In particular, we use $MA_{geo,i} = \sum_{s \neq i} \frac{1}{\mathbf{E}[T(i,s)]^{\sigma-1}}$ as a measure of market access for city i, where $\mathbf{E}\left[T\left(i,s\right)\right]$ is the expected pre-containerization trade cost between cities i and s. All our results are robust to using more complex measures of market access that include the populations of other cities, s.

[79]Table 8 also shows that the share of port cities and country population have a significant positive effect on country-level gains. These variables are positively correlated with GDP per capita (with correlation coefficients of 0.18 and 0.17, respectively). Hence, they work against poor countries gaining more from containerization.

[80]In Appendix Figure A.15, we show the association between the welfare gains and market access conditional on the covariates in column (6) of Table 8. The results are qualitatively similar if we examine the unconditional association.

In summary, poor countries gained more from containerization than rich ones. This section has shown that this is driven in part by the novel forces our model includes. Two significant drivers behind the relationship are productivity in non-port activities and pre-containerization market access. The first drives endogenous specialization based on comparative advantage. Second, countries with worse initial market access gained more from containerization, partly due to lower shipping costs and partly as a result of endogenous port development.

8 Conclusion

The containerization shock studied in this paper allows us to shed light on the economic effects of port development. Much like other transport infrastructure improvements, at the local level, port development tends to make a location attractive for firms and workers through the standard market access effect. However, different to many transport infrastructure improvements such as roads and railways, the fact that ports occupy vast amounts of space in their host cities also leads to a strong opposing force that tends to crowd out population. The paper has shown that in the case of the containerization shock, this endogenous crowding out force is strong and has the potential to matter for both the local and aggregate economic effects of port development.

Though the analysis in this paper is positive, it offers some tentative implications for where port development is likely to have the biggest beneficial impact. On the one hand, the recent aggressive port development strategy followed by some developing country cities such as Colombo, Sri Lanka seems promising. These are cities where the opportunity cost of land remains relatively low given the low productivity of non-port activities. Our results suggest that port development could lead to relatively large benefits for the entire country. On the other hand, our findings cast some doubt on the wisdom of further developing or maintaining high levels of port activity in some of the world's most expensive cities such as Hong-Kong and Singapore. While these cities arguably benefited enormously from their position as important ports historically (at a time when they were also far poorer relative to the rest of the world), subsequent productivity growth *outside* the port sector has made the opportunity cost of the land occupied by the port extremely high. Our findings suggest that the 'Hong-Kongs' and 'Singapores' of the world may benefit from following the path of cities such as London (United Kingdom) – a city at the center of world trade for many decades, but one that now houses Canary Wharf, an important second financial district, on redeveloped land once occupied by the port.

References

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

In terms of statistical significance, the coefficient on market access is significant at a 10% level in our baseline model (see column (6) of Table 8), but insignificant in both Benchmark 1 and Benchmark 2.

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

Allen, T. and C. Arkolakis (2014). Trade and the Topography of the Spatial Economy. *Quarterly Journal of Economics 129*(3), 1085–1140.

Allen, T. and C. Arkolakis (2019). The Welfare Effects of Transportation Infrastructure Improvements.

Allen, T. and D. Atkin (2016). Volatility and the Gains from Trade.

Altomonte, C., I. Colantone, and L. Bonacorsi (2018). Trade and Growth in the Age of Global Value Chains. Technical report, BAFFI CAREFIN Centre Research Paper.

Anderson, J. (1979). A Theoretical Foundation for the Gravity Equation. *American Economic Review 69*(1), 106–116.

Armenter, R., M. Koren, and D. Nagy (2014). Bridges.

Bernard, A., J. Eaton, J. Jensen, and S. Kortum (2003). Plants and Productivity in International Trade. *American Economic Review 93*(4), 1268–1290.

Bernhofen, D., Z. El-Sahli, and R. Kneller (2016). Estimating the Effects of the Container Revolution on World Trade. *Journal of International Economics 98*, 36–50.

Bleakley, H. and J. Lin (2012). Portage and Path Dependence. *Quarterly Journal of Economics 127*, 587–644.

Blonigen, B. and W. Wilson (2008). Port Efficiency and Trade Flows. *Review of International Economics 16*(1), 21–36.

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2020). Geography, transportation, and endogenous trade costs. *Econometrica 88*(2), 657–691.

Brooks, L., N. Gendron-Carrier, and G. Rua (2019). The Local Impact of Containerization.

Campante, F. and D. Yanagizawa-Drott (2018). Long-range Growth: Economic Development in the Global Network of Air Links. *Quarterly Journal of Economics 133*(3), 1395–1458.

Ciccone, A. and R. Hall (1993). Productivity and the Density of Economic Activity. *National Bureau of Economic Research* (Working Paper 4313).

Coşar, A. and P. Fajgelbaum (2016). Internal Geography, International Trade, and Regional Specialization. *American Economic Journal: Microeconomics 1*(8), 24–56.

Conley, T. (1999). Gmm Estimation with Cross Sectional Dependence. *Journal of Econometrics 92*(1), 1–45.

Corbett, M. (2010). *The History and Transformation of the Port of San Francisco, 1848-2010*. San Francisco Architectural Heritage.

Corden, W. and J. Neary (1982). Booming sector and de-industrialisation in a small open economy. *Economic Journal 92*(368), 825–848.

Coşar, A. and B. Demir (2018). Shipping Inside the Box: Containerization and Trade. *Journal of International Economics 114*, 331–345.

Desmet, K. and J. Rappaport (2017). The Settlement of the United States, 1800–2000: The Long Transition Towards Gibrat's Law. *Journal of Urban Economics 98*, 50–68.

Donaldson, D. and R. Hornbeck (2016). Railroads and American Economic Growth: A "market access" Approach. *Quarterly Journal of Economics 131*(2), 799–858.

Donaldson, D. and A. Storeygard (2016). The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives 30*(4), 171–98.

Ducruet, C., S. Cuyala, and A. E. Hosni (2018). Maritime Networks as Systems of Cities: The

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

Duranton, G. and D. Puga (2019). Urban Growth and its Aggregate Implications.

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

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

Fajgelbaum, P. and S. Redding (2018). Trade, Structural Transformation and Development: Evidence from Argentina, 1869-1914.

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

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

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

Ganapati, S., W. F. Wong, and O. Ziv (2020). Entrepôt: Hubs, Scale, and Trade Costs.

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

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

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

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

Heiland, I., A. Moxnes, K. H. Ulltveit-Moe, and Y. Zi (2019). Trade From Space: Shipping Networks and The Global Implications of Local Shocks.

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

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

Jedwab, R. and A. Storeygard (2020). The Average and Heterogeneous Effects of Transportation Investments: Evidence from Sub-Saharan Africa 1960-2010. *National Bureau of Economic Reserach* (Working Paper 27670).

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

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

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

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

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

Maurer, S. and F. Rauch (2020). Economic Geography Aspects of the Panama Canal.

McKinsey & Company (1972). Containerization: A Five-Year Balance Sheet. Technical report, McKinsey & Company.

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

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

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

Nagy, D. (2018). Trade and Urbanization: Evidence from Hungary.

Nagy, D. (2020). Hinterlands, City Formation and Growth: Evidence from the U.S. Westward Expansion.

New York Port Authority (1958). 1958 Annual Report. Technical report, New York Port Authority.

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

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

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

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

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

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

Riffenburgh, R. (2012). A Project History of the Port of Long Beach 1970 to 2010. Technical report, Port of Long Beach.

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

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

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

Santamaría, M. (2020). The Gains from Reshaping Infrastructure: Evidence from the division of Germany.

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

UNCTAD (1971). Review of Maritime Transport. Technical report, United Nations Publication.

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

Wong, W. F. (2017). The Round Trip Effect: Endogenous Transport Costs and International Trade.

A Tables

Table 1: Depth predicts shipping flows, but only after 1960

	Dependent Variable: ln(Shipment)				
Independent Variables	(1)	(2)	(3)	(4)	(5)
Depth × post 1970					0.247***
					(0.059)
					{0.052}
Depth × 1960	-0.051	0.029	0.050	-0.055	
	(0.063)	(0.069)	(0.066)	(0.068)	
Depth × 1970	0.222***	0.233***	0.278***	0.213***	
	(0.069)	(0.077)	(0.082)	(0.071)	
Depth × 1980	0.188**	0.212**	0.291***	0.192**	
	(0.079)	(0.085)	(0.090)	(0.081)	
Depth × 1990	0.255***	0.222**	0.312***	0.283***	
	(0.086)	(0.087)	(0.099)	(0.087)	
Observations	2765	2765	2765	2360	2765
R-squared	0.126	0.248	0.131	0.142	0.126
Number of cities	553	553	553	472	553
Year FE	✓	✓	✓	✓	✓
City FE	✓	✓	✓	✓	✓
Population 1950 × Year	✓	✓	✓	✓	✓
Coastline × Year FE	×	✓	×	×	×
Saiz × 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. Standard errors clustered at the city level in parentheses, Conley standard errors to adjust for spatial correlation in curly brackets. *** $p<0.01$, ** $p<0.05$, * $p<0.1$ (significance refers to clustered standard errors).

Table 2: The local causal effect of shipping on population is not distinguishable from zero

	Panel regression						Long difference			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Indep. Variables	ln(Pop)	ln(Pop)	ln(Ship)	ln(Pop)	ln(Ship)	ln(Pop)	Δln(Pop)	Δln(Pop)	Δ ln(Ship)	Δln(Pop)
ln(Shipment)	0.013***	0.015								
	*0.030****	*0.035*								
	(0.005)	(0.049)								
Δ ln(Shipment)							0.013	0.006		
							0.052	*0.022*		
							(0.009)	(0.073)		
Depth									0.272***	0.002
									*0.134****	*0.003*
									(0.086)	(0.020)
Depth $\times$ post 1970			0.268***	0.004						
			*0.143****	*0.005*						
			(0.058)	(0.013)						
Depth $\times$ 1960					-0.042	-0.003				
					(0.064)	(0.008)				
Depth $\times$ 1970					0.246***	0.007				
					(0.069)	(0.013)				
Depth $\times$ 1980					0.213***	-0.002				
					(0.079)	(0.017)				
Depth $\times$ 1990					0.280***	0.002				
					(0.086)	(0.020)				
Observations	2734	2734	2734	2734	2734	2734	531	531	531	531
Number of cities	552	552	552	552	552	552				
Year FE	✓	✓	✓	✓	✓	✓	×	×	×	×
City FE	✓	✓	✓	✓	✓	✓	×	×	×	×
Population 1950 $\times$ Year	✓	✓	✓	✓	✓	✓	×	×	×	×
Population 1950	×	×	×	×	×	×	✓	✓	✓	✓
Specification	OLS	2SLS	FS	RF	dyn FS	dyn RF	OLS	2SLS	FS	RF
KP F-stat		21.13						9.98		

Notes: "Depth" indicates the port suitability measure. It is interacted with decade dummies or indicator variables for decades including and after 1970, as indicated. Standardized coefficients in italics underneath the baseline coefficients. Notation for specification as follows: 'FS' refers to the first stage, 'RF' to the reduced form, 'dyn FS' to the fully flexible first stage and 'dyn RF' to the fully flexible reduced form. Standard errors clustered at the city level (Appendix Table A.8 reports Conley standard errors for the main results). *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Containerization increased shipping more in low rent cities

	ln(Shipment)	
Independent Variables	(1)	(2)
Depth × post 1970	0.464***	0.566***
	(0.138)	(0.152)
	{0.094}	{0.118}
Depth × Saiz × post 1970	-0.408*	-0.707**
	(0.220)	(0.323)
	{0.153}	{0.237}
Saiz × post 1970		0.975
		(0.804)
		{0.588}
Observations	2765	2765
R-squared	0.128	0.129
Number of cities	553	553
Year FE	✓	✓
City FE	✓	✓
Population 1950 × Year	✓	✓

Notes: "Depth" indicates the port suitability measure. "Saiz" is the Saiz land rent proxy defined in Saiz (2010). Each measure is interacted with an indicator for decades including and after 1970, and we also include the triple interaction term in the regression, which is the coefficient of interest. Standard errors clustered at the city level in parentheses, Conley standard errors to adjust for spatial correlation in curly brackets. *** $p<0.01$, ** $p<0.05$, * $p<0.1$ (significance refers to clustered standard errors).

Table 4: Model-inspired specification: Disentangling market access effect and crowding out effect

Independent Variables	(1) ln(Population)	(2) ln(Population)	(3) ln(Population)	(4) ln(Population)	(5) ln(Shipment)	(6) ln(Market Access)
ln(Shipment)	0.015***	0.014	-0.001	-0.159**		
	(0.005)	(0.048)	(0.006)	(0.065)		
	{0.005}	{0.038}	{0.005}	{0.051}		
ln(Market Access)			1.512***	7.103***		
			(0.536)	(0.795)		
			{0.317}	{0.854}		
Depth × post 1970					0.275***	0.007***
					(0.058)	(0.001)
					{0.051}	{0.001}
Market Access IV					7.188	1.927***
					(5.428)	(0.140)
					{5.748}	{0.188}
Observations	2696	2696	2696	2696	2696	2696
R-squared	0.718	0.718	0.735	0.417		
Number of cities	544	544	544	544	544	544
Year FE	✓	✓	✓	✓	✓	✓
City FE	✓	✓	✓	✓	✓	✓
Population 1950 × Year	✓	✓	✓	✓	✓	✓
Specification	OLS	2SLS	OLS	2SLS	FS	FS
KP F-stat		22.07		9.63		

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

Table 5: Calibration of structural parameters

Parameter	Target
$\alpha = 0.06$	Agglomeration externalities (Ciccone and Hall, 1993)
$\gamma = 0.84$	Non-land share in production (Desmet and Rappaport, 2017)
$\eta = 0.15$	Migration elasticity (Kennan and Walker, 2011)
$\sigma = 4$	Elasticity of substitution across tradables (Bernard et al., 2003)
$\theta = 203$	Idiosyncratic shipping cost dispersion (Allen and Arkolakis, 2019)
$\lambda = 0.074$	Congestion externalities in ports (Abe and Wilson, 2009)

Table 6: The causal effect of shipping on local population in the data, 'baseline model', and 'benchmark models'

	Δln(Population)			
	Data	Model	Benchmark 1	Benchmark 2
Independent Variables	(1)	(2)	(3)	(4)
Δln(Shipment)	0.006	0.001	0.015**	0.018***
	0.022	*0.006*	*0.124***	*0.140****
	(0.073)	(0.007)	(0.006)	(0.007)
Observations	531	553	553	553
Specification	2SLS	2SLS	2SLS	2SLS
KP F-stat	9.98	595.88	666.45	662.22

Notes: Column (1) uses depth as IV for shipping, controlling for population in 1950, which is equivalent to using residualized depth as an IV. Columns (2) to (4) use residualized depth as IV, which is the variation that we feed into the model to simulate the counterfactual. Standardized coefficients in italics underneath the baseline coefficients. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 7: The aggregate welfare effects of containerization - Sensitivity analysis

Model	Welfare effect (%)
1. Baseline	3.95
2. 20% higher β in inversion	3.98
3. 20% lower β in inversion	3.94
4. Counterfactual β implying 65% increase in port share of New Orleans	3.97
5. Counterfactual β implying 85% increase in port share of New Orleans	3.93
6. No depth-dependent change in $\nu(r)$	4.30
7. Larger ν_{CF}: implies 30% change in total transshipment costs	4.60
8. Smaller ν_{CF}: implies 20% change in total transshipment costs	3.26
9. Monopolistic competition	4.34

Table 8: The determinants of the country-level welfare gains from containerization

	Change in country-level welfare					
Independent Variables	(1)	(2)	(3)	(4)	(5)	(6)
ln(GDP per capita)	-0.161**	-0.288***	-0.293***	-0.390*	0.216	
	(0.068)	(0.061)	(0.065)	(0.235)	(0.383)	
Share of port cities		0.693***	0.699***	0.699***	0.727***	0.723***
		(0.071)	(0.064)	(0.064)	(0.072)	(0.070)
ln(Population)			0.015	-0.008	0.206*	0.144***
			(0.051)	(0.085)	(0.121)	(0.055)
ln(Average Productivity)				0.107	-0.575	-0.340***
				(0.264)	(0.432)	(0.071)
ln(Average MA_{geo})					-0.199	-0.156*
					(0.137)	(0.086)
Observations	167	167	167	167	167	167
R-squared	0.026	0.490	0.490	0.491	0.504	0.503

Notes: The table shows standardized beta coefficients. The change in welfare is measured in %. All independent variables correspond to their counterfactual values. Country averages weighted by counterfactual city populations. $MA_{geo,i} = \sum_{s \neq i} \frac{1}{\mathbf{E}[T(i,s)]^{\sigma-1}}$. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B Figures

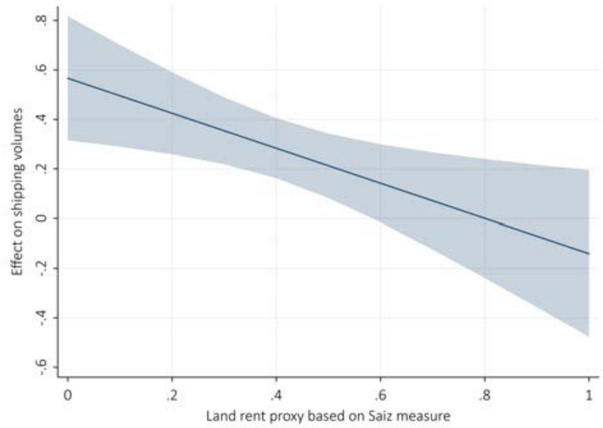

Figure 1: The estimated effect of depth on shipping evaluated at different values of the Saiz land rent proxy in the data

Notes: The figure shows the estimated γ coefficient from equation (2) evaluated at different values of the Saiz land rent proxy. The corresponding regression results are reported in column (2) of Table 3.

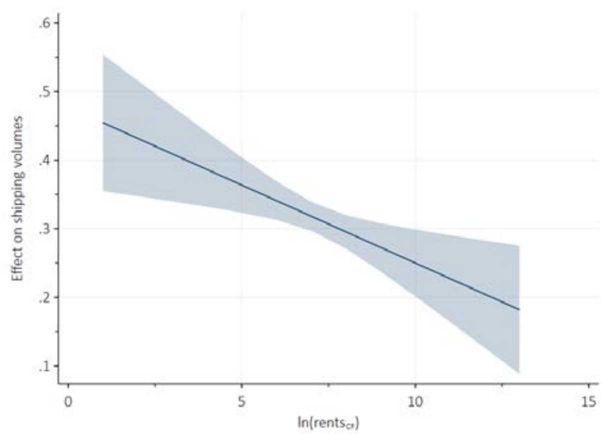

Figure 2: The estimated effect of depth on shipping evaluated at different values of the counterfactual land rents in the model

Notes: The figure shows the estimated γ coefficient from equation (14) evaluated at different values of the counterfactual land rents. The corresponding regression results are reported in Appendix Table A.20.