

All aboard: The aggregate effects of port development

César Ducruet¹, Réka Juhász², Dávid Krisztián Nagy³, and Claudia Steinwender⁴

¹CNRS

²Columbia University

³CREI, Universitat Pompeu Fabra and Barcelona GSE

⁴MIT Sloan *

July 3, 2019

Abstract

This paper studies the distributional and aggregate economic effects of new port technologies developed in the second half of the 20th century. We show that new technologies have led to a significant reallocation of shipping activity from large to small cities. This was driven by a land price mechanism; as new port technologies are more land-intensive, ports moved from large, high land price cities to smaller, lower land price ones. We add endogenous port development to a standard quantitative model of cross-city trade to account for both the benefits and the costs of port development. According to the model, the adoption of new port technologies leads to benefits through increasing market access but is costly, requiring the extensive use of land, suggesting a reallocation of shipping activities towards cities with low land prices and thus net gains from new port technologies that are heterogeneous across cities. Counterfactual results suggest that new port technologies led to sizable aggregate gains for the world economy, with substantial heterogeneity in the effects across countries. More generally, accounting for the costs of port infrastructure development endogenously has the potential to alter the size and distribution of the gains from trade.

*We thank David Atkin, Don Davis, Dave Donaldson, Joseph Doyle, Matt Grant, Gordon Hanson, Tom Holmes, Amit Khandelwal, Jim Rauch, Roberto Rigobon, Esteban Rossi-Hansberg and Tavneet Suri for helpful comments and discussions. We thank Bruce Blonigen for kindly sharing data. Naman Garg, Yi Jie Gwee, Rodrigo Martínez Mazza, Emanuela Migliaccio and Howard Zhiao Zhang provided outstanding research assistance. Réka acknowledges funding from the Provost's Grant for Junior Faculty who Contribute to the Diversity Goals of the University.

Introduction

Cities across the income distribution spend heavily on port development.¹ There is widespread belief that the gains to developing port infrastructure are high.² However, port services, and in particular transshipment services, are a low value added economic activity, suggesting that port development must exert its positive influence through linkages to other parts of the economy (OECD, 2014). Spatial theory provides at least one such linkage; the well-known market access effect. By reducing transport costs, port development makes a location attractive for firms and consumers alike. Indeed, a recent literature has emerged showing that trade induced city development is important for both aggregate development and structural transformation (Bleakley and Lin, 2012; Armenter et al., 2014; Coşar and Fajgelbaum, 2016; Nagy, 2017; Fajgelbaum and Redding, 2018).

However, alongside the benefits, port activities can also entail significant costs for a city. Modern ports have become increasingly land-intensive to the point where large ports today occupy an enormous share of a city's available land supply. According to some estimates, the ports of Antwerpen and Rotterdam occupy more than 30% of the metropolitan area, a staggering number (OECD, 2014). Through their land usage, modern ports may impose significant costs on a city. In recent years, concerns have been raised that port development may constrain city development through this mechanism.³ As land prices increase systematically with city size (Combes et al., 2018), the opportunity cost of port development may vary across cities.

In this paper, we examine the aggregate and distributional effects of port development induced by the arrival of new port technologies in the 1960s, allowing for both the benefits and the costs that these new technologies may entail. By reducing trade costs, new port technologies, particularly containerization, have been found to account for a sizable proportion of the increase in trade flows during the second wave of globalization (Bernhofen et al., 2016;

¹Two prominent recent examples from very different parts of the income distribution are Colombo and Hong Kong. Colombo's port development is described in Herrera Dappe and Suárez-Alemán (2016), while recent media reports describing plans to build housing units above the container port illustrate the lengths city officials are willing to go to maintain Hong-Kong's status as a premier port amidst a drastic housing shortage (South China Morning Post, retrieved July 1, 2019).

²For example, developers of Colombo's port hope to transform the city into "(...) something more than just a place to transship containers. It is to be a major financial center, rivaling Singapore to the east and Dubai to the west (...)" (Forbes Magazine, 2016, retrieved July 1, 2019). A recent report by the OECD summarizes more generally the purported link between port development and urban growth (OECD, 2014).

³For example, according to the OECD, "Agglomeration effects and high job density are generally considered to be factors of urban economic growth and these agglomeration effects may be constrained by the presence of large port areas." (OECD, 2014, p.51)

Coşar et al., 2018; Levinson, 2010). However, differently to other transportation technology improvements such as the railroads, little is known about their aggregate and distributional effect on economic activity.⁴

Studying the effects of new port technologies is particularly interesting, as differently to other transport infrastructure such as roads or railways, modern ports place a significant strain on locally scarce land supply in cities. This raises the possibility of large, heterogeneous costs of developing port infrastructure, suggesting that not all cities are equally well-suited to hosting modern ports and that new port technologies may significantly alter the spatial distribution of economic activity. Indeed, new port technologies have been argued to have substantially changed the economic geography of cities. For example, in an influential book, Levinson (2010) argues that “Cities that had been centers of maritime commerce for centuries, such as New York and Liverpool, saw their waterfronts decline with startling speed [...] Sleepy harbors such as Busan and Seattle moved into the front ranks of the world’s ports.”

We use a novel and unique dataset on bilateral shipping flows across ports worldwide for the period 1950-1990 to provide what is to the best of our knowledge the first quantitative, causal estimates of the reallocation of ports across cities as a consequence of these new technologies. We provide evidence for the mechanism that we argue drives these changes and build a quantitative spatial model to estimate their effect on economic activity.

In the first part of the paper, we estimate the extent to which the reallocation of port activity from big to small cities, as argued by Levinson (2010), is indeed a systematic feature of the data and whether it is caused by the arrival of new port technologies. We confirm both. We find that starting in the 1960s, ports reallocated from the largest cities to smaller ones. We show that this effect was driven by initially smaller cities increasing their shipping activities disproportionately relative to large ones. Next, we show that part of this reallocation was causal. We isolate exogenous variation in the suitability of cities to new port technologies by developing an area-based measure of natural depth, building on Brooks et al. (2018). Using this source of variation, we show that cities exogenously better suited to new port technologies increased their shipping volumes disproportionately after 1970, but not before. In line with the overall reallocation of ports from big to small cities, we find that our measure of exogenous port suitability predicts disproportionately larger increases in shipping flows for cities that were initially smaller. This novel stylized fact motivates the

⁴For example, Donaldson (2018); Donaldson and Hornbeck (2016) study the economic effect of the railroads in a quantitative trade model.

remainder of the paper.

We next turn to identifying the mechanisms that account for this reallocation. Our analysis is informed by two aspects of the new technology. First, we document that new technologies dramatically reduced transshipment costs at the port. This is in contrast to other transportation technology improvements such as the steamship or the railway that predominantly reduced distance based transport costs. Second, we document that these technologies led to ports becoming far more land-intensive. This was partly because containerized shipping requires more space during loading, unloading and transshipment, and partly because modern ports operate at a larger scale. Informed by these aspects of the empirical setting, we hypothesize that the reallocation of shipping was driven by a land-price mechanism. As new port technologies meant more land was required for port development, hosting ports in large, high land price cities became more expensive, making the reallocation of ports to smaller cities more attractive. We show three pieces of empirical evidence in support of this mechanism.

First, we document that cities exogenously better suited to becoming a modern port increased their shipping disproportionately after the 1970s, but more so in places where land supply was elastic. That is, cities equally well-suited to hosting modern ports in terms of their depth measure increased their shipping flows to a larger extent if there was more land available for the city to expand. Second, we show that larger (more populous) cities dedicated a *smaller* share of land to the port after the arrival of new port technologies. This is despite the fact that larger cities had systematically *higher* shipping flows, as predicted by gravity models. Moreover, this specialization away from port based activities was not present in big cities prior to the arrival of new port technologies. Before their arrival, larger cities allocated a larger share of their land to the port than smaller ones. Third, we estimate the casual effect of shipping on population in port cities and in nearby inland cities. We find that the net effect of shipping on population is positive in port cities, however it is somewhat smaller than the effect on cities that are located inland, but are close to a port. This suggests that nearby inland cities enjoy similar benefits of the port as the port cities, but do not face the same costs.

To rationalize these facts and develop a tool for quantitative general equilibrium analysis, the second part of the paper builds a model of trading cities. In the model, land owners in port cities can use part of their land to provide transshipment services to firms shipping through the port. Land owners can lower the cost of transshipment by *developing the port*, that is, by increasing the share of land dedicated to transshipment. New port technologies,

by being more land-intensive, make this relationship between land use and transshipment costs more pronounced than before. Endogenous transshipment costs enter overall shipping costs in a way that the model preserves the gravity equation of trade flows. The model provides a rich set of predictions on the extent to which cities develop their port. On the one hand, cities develop their port more if they face a large amount of shipping through the port, implying larger benefits from port development. On the other hand, cities develop their port less if they face high land rents, implying higher costs of port development.

The model also provides testable predictions on two opposing effects of port development on city population. On the one hand, port development benefits the city as it lowers shipping costs, hence increases the city's access to trade, an effect that we call the *market access effect*. Everything else fixed, the market access effect draws people into the city. On the other hand, port development is costly as it requires land. Developing the port involves reallocating land from other purposes, thus making the city a worse place to live, an effect that we call the *crowding-out effect*. Everything else fixed, the crowding-out effect pushes people out of the city. Whether a city ultimately gains in population is the outcome of the trade-off between these two forces. Differences across cities in geographic location and land rents imply that the trade-off between the two forces plays out differently for different cities, making them benefit differentially from the arrival of new port technologies and thus triggering a reallocation of economic activity across cities.

Informed by the model's predictions on city population driven by the costs and benefits of port development, we re-estimate the causal effect of increased shipping flows on population, controlling for market access as guided by the model. Our causal estimates point to a negative effect of shipping on city population once market access is controlled for. In line with the predictions of the model, once the indirect (linkage) effect of increased shipping is accounted for through the market access channel, the direct effect of shipping on population is negative, consistent with a crowding-out effect of port activities.

In the final part of the paper, we take the model to the data in order to quantify the aggregate and distributional effects of new port technologies. We use data on shipping flows, city GDP and population in 1990 to back out the city-specific fundamentals that rationalize the data and allow us to conduct counterfactuals. To illustrate the quantitative effects of new port technologies, we conduct a counterfactual in which we change the parameters of the endogenous transshipment cost function in a way that mimics transshipment technology in the 1950s. Importantly, we show that moving from the old to new port technologies leads to a reallocation of shipping from large to small cities consistent with our motivating stylized

fact. In this counterfactual, shipping flows increase by 92.6%, while aggregate welfare and real GDP increase by 5.3% and 5.7%, respectively. Moreover, the gains in welfare vary substantially across countries, ranging from zero to as much as 50%.

While we think of these numbers as illustrative of the mechanisms in the model as opposed to being precise estimates of the effect of new port technologies, we believe there are important conclusions to be drawn from this preliminary evidence. Compared to other estimates of the gains from increased trade, the aggregate and distributional effects in our setting seem to be relatively sizable. In our view, this underscores the importance of thinking carefully about the costs of infrastructure development. In particular, one interpretation of the sizable aggregate gains found in this paper is that new port technologies led to large reallocation of economic activity in a way that compounds welfare effects through population movements. As ports reallocated from large, high-rent, high-productivity cities to smaller, lower rent, lower productivity cities, they were able to reap the gains from increased specialization. Larger cities became more specialized in their comparative advantage sector, non-port activities. More generally, this suggests that accounting for the costs of infrastructure in an endogenous way has the potential to deepen our understanding of the gains from trade and how they are distributed.

The paper is organized as follows. Section 1 describes new port technologies. In Section 2, we introduce the most important data sources used in the paper. Section 3 presents the reduced form effects of new port technologies. Section 4 discusses the model. Section 5 revisits the empirical evidence in light of the model predictions. Section 6 describes how we take the model to the data, while Section 7 assesses the model fit. Finally, Section 8 discusses the preliminary counterfactual results.

Related literature

A recent, growing literature has shown evidence that access to trade leads to local benefits inducing city development (Bleakley and Lin, 2012; Armenter et al., 2014; Coşar and Fajgelbaum, 2016; Nagy, 2017; Fajgelbaum and Redding, 2018). We contribute to this literature by showing that trade induced development can also have substantial, heterogeneous local costs. The crowding out mechanism that drives the cost side in our setting relates the paper to the “Dutch-disease” literature which shows that booming industries lead to significant costs by increasing local factor prices and crowding out other (tradeables) sectors (Corden and Neary, 1982; Krugman, 1987; Allcott and Keniston, 2017). Relative to this literature, our setting contains the potential for gains as well as costs, through the input-output linkages to other sectors in the form of market access. One contribution of this paper

is to generalize the predictions from these two, seemingly disparate literatures which have focused on either the cost or benefit side of booming industries. When estimating the effect of a booming sector that competes for locally scarce resources, it is important to account for potential benefits through input-output linkages between the booming sectors and the rest of the economy. At the same time, it is also important to account for the costs in terms of potential crowding out through competition for locally scarce resources.

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 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 (2017) and Santamaría (2018), who consider endogenous road construction in multi-location models of economic geography. Unlike them, we focus on endogenous port development as opposed to road infrastructure, and solve for the decentralized equilibrium as opposed to the optimal allocation to improve our understanding of the large-scale changes in shipping and economic activity observed in the data.

Finally, our paper is related to a large literature studying the effects of transport infrastructure improvements.⁵ There is a growing literature studying the effects of new port technologies, in particular the effects of containerization (Hummels, 2007; Bernhofen et al., 2016; Coşar et al., 2018; Holmes and Singer, 2018; Brooks et al., 2018; Gomtsyan, 2016, Altomonte et al. 2018).⁶ Most closely related is Brooks et al. (2018) who study the effect of containerization on local economic outcomes across US counties. Our main contribution to this literature is twofold. First, motivated by our novel stylized fact documenting the reallocation of ports from big to small cities, this paper focuses on understanding how the land-mechanism drives heterogeneous costs of port development across cities. Second, to the best of our knowledge, this is the first paper seeking to quantify the aggregate effects of new port technologies through the lens of a quantitative economic geography model.

⁵Redding and Turner (2015) provides an overview of recent developments in this literature.

⁶More broadly, the paper also relates to Brancaccio et al. (2017) who endogenize trade costs in the shipping sector in order to study its effect on trade. Relative to that paper, the focus in this paper is on understanding how the *location* of modern ports affects the spatial distribution of economic activity.

1 New port technologies

The middle of the 20th century witnessed a number of revolutionary changes to shipping, and in particular, port technology. The best-known and perhaps most important of these was the development and adoption of containerized shipping worldwide, that is the handling of cargo in standardized boxes (called containers). The nature of handling large containers heralded a host of other innovations in port technologies. These included automation, in particular the introduction of large cranes for loading and unloading cargo, ship sizes that are 17 times larger than in the 1950s (Brooks et al., 2018) and computerization, which for example facilitates the optimal unloading of cargo taking into account the weight distribution in the ship (Levinson, 2010). The use of these technologies is not limited to containerized cargo handling, but instead a general feature of modern ports. Henceforth, we refer to this bundle of new innovations as *new port technologies*. We believe this distinction is important, as while containerization is likely to be the major driver of cost reductions and changes in port technologies, our focus will not be limited to understanding containerization per se, but rather in estimating the effects of all new port technologies. This is key, as both our data and our reduced form empirical strategy are designed and suited to capturing the effect of new port technologies that are not limited to containerization.

The idea behind containerization was remarkably simple: instead of handling discrete cargo items individually, why not put cargo into boxes and handle those? Containerized shipping of cargo was initially introduced on domestic routes between US ports, but the technology was rapidly adopted and importantly, standardized worldwide in the 1960s (Rua, 2014). By the mid-2000s, containerized shipping was estimated to account for 70% of the volume of general cargo (excluding oil, fertilizers, ore and grain) (Rua, 2014). Similarly to other transportation technology improvements, containerization and modern port technologies more generally reduced transportation costs and increased trade flows (Hummels, 2007; Bernhofen et al., 2016; Coşar et al., 2018). There are two aspects of the nature of these new technologies that are particularly important to the question examined in this paper. In the following, we discuss these in more detail.

1.1 Modern ports take up a lot of land

First, new port technologies require far more land than traditional ones in order to accommodate the cranes that move the containers and for the marshalling of containers and trucks (Rua, 2014). Modern ports also need a larger scale to operate efficiently given their increased

capital and land-intensity (Levinson, 2010; OECD, 2014). Taken together, this implies that modern ports have the potential to occupy large areas of the city. Indeed, the OECD estimates that the port of Antwerpen and Rotterdam occupy more than 30% of the metropolitan area of the city, a staggering share. Even in larger, extremely high land-price cities such as Hong-Kong, the port covers an estimated 5% of the city (OECD, 2014).

While we are aware of no readily available measures of port size across time for our worldwide sample of cities, given the central role played by the local land market in this paper, it is important to understand whether port sizes have indeed increased over time. To this end, we collected wharf length data for ports around the world for one year each decade between 1950-1990.⁷ Figure 1 shows the distribution of port length throughout the decades. There is very little change in the distribution between 1950 and 1960, and a distinct and steady rightward shift of the distribution in the following decades, consistent with the timing of the development of new technologies in shipping. While this is a highly imperfect measure of port area, it does confirm for our worldwide sample that port area seems to have become systematically larger after the introduction of new port technologies.⁸

1.2 Modern port technologies reduced transshipment costs

Second, containerization reduced costs most dramatically in transshipment. As Krugman (2011) writes, “The ability to ship things long distances fairly cheaply has been there since the steamship and the railroad. What was the big bottleneck was getting things on and off the ships. A large part of the costs of international trade was taking the cargo off the ship, sorting it out, and dealing with the pilferage that always took place along the way. So, the first big thing that changed was the introduction of the container.” Today, large container ports that house acres upon acres of orderly stacked standardized boxes are such an ubiquitous feature of cities that it is hard to imagine that only a few decades ago, ports used very different technologies. Breakbulk shipping, the technology containerization displaced, was a complicated and costly method of transshipping freight at the port. Each item of cargo was handled individually making the loading and unloading of ships expensive both in terms of freight costs and in terms of time a ship spent in port. According to Bernhofen et al. (2016), prior to containerization, often two-thirds of a ships’ time would be spent in port. Industry experts estimated that using breakbulk shipping technology,

⁷The port length data is discussed in more detail in Section 2.

⁸The most important drawback of our data is that it fails to capture the increased land usage of a port due to space dedicated to cranes and the storage of containers. However, this implies that we are probably vastly underestimating the increase in port size over this time period.

the handling of cargo at the port accounted for the lion’s share of freight costs (Levinson, 2010); transshipment costs were estimated to account for 49% of the total transport cost on shipment from the US to Europe (Eyre, 1964). Moreover, containerization also made intermodal transport vastly more efficient as the same standardized boxes were adopted for rail freight and trucking, further facilitating the transshipment of cargo. Given the central importance attributed to transshipment cost reductions by the literature, we model new port technologies as a reduction in transshipment costs in the model.

2 Data

Our analysis builds on a number of novel data sources.⁹ Crucial to our analysis is a dataset of worldwide bilateral ship movements for the years 1950-1990 at the port level.¹⁰ These data were extracted for one week samples for each decade from the Lloyd’s List. Differently to other port level data sources, as these ship movements capture both containerized and non-containerized shipping flows, we are able to compare the economic geography of shipping and its effect on cities both before and after the arrival of new port technologies. We know of no other data source that has a similar coverage of shipping flows across time and space, especially at such a detailed port-level disaggregation. An important limitation, however, is that we cannot 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.

As we are interested in the economic effects of new port technologies on cities, we match the shipping data to city population. We use data on city population for locations with greater than 100,000 inhabitants from “Villes Géopolis” for the years 1950-1990 (Geopolis city, henceforth) (Moriconi-Ebrard, 1994).¹¹ Ports from the shipping data were hand-matched to cities based on whether they were located within the urban agglomeration of a city in the Geopolis dataset, allowing for multiple ports to be assigned to one city. We define port cities in a time invariant manner; a port city with positive shipping flows in

⁹In this section, we only discuss the most important data sources. All remaining data are discussed in the Online Appendix.

¹⁰The data are discussed in more detail in Ducruet et al. (2018).

¹¹The advantage of these data relative to sources such as the 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. 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 strictly fewer than 100,000 inhabitants.

only one period will be classified as a port city for all years. Table 1 contains summary statistics on our matched shipping-city sample. Of the 2,636 cities in the Geopolis dataset, 553 have at least one port assigned, we label them port cities. Besides the port cities, we have information for 1,592 ports that are not assigned to a city in the Geopolis dataset. Our empirical analysis focuses predominantly on estimating the effects of new port technologies on the 553 port cities, though we also show some effects on nearby Geopolis cities. As we are interested in estimating the general equilibrium effects of port technologies, our quantitative estimation covers the full set of 2,636 Geopolis cities (port and non-port cities). Figure 2 plots the port and non-port cities.

Figure 3 visualizes the shipping network at the beginning and the end of our sample period.¹² The nodes in this network are the port cities and their size is proportionate to the population of the city in each period. The width of the edges are proportionate to the size of shipping between two nodes. A cursory look at the data anticipates the results we will show formally in the next section; between 1950 and 1990, shipping activity reallocated from the largest cities in the world in terms of population to much smaller ones.

Finally, the size of the port and the amount of land it occupies as a share of the city is central to our analysis. We are aware of no readily available measures for these across time and covering our sample of port cities. In their absence, we have constructed an imperfect approximation using the best available information. We have collected data on the length of wharves (“port length”) for port cities from the annual publication “Ports of the World”. This is available for 83% of the ports in our sample for at least some years. For a small subset of ports (14), we also observe the total area of land occupied by the port in 1990 from the same publication. We use the scaling factor between port length and total port area on this subsample to transform port length into port area on the remaining sample.

We measure the area of the city using nightlight satellite data to identify the boundary of the city in 1990.¹³ The share of the city occupied by the port is the estimated port area divided by the total estimated area of the city. While this is an admittedly highly imperfect measure with multiple sources of measurement error, we find that it shows patterns that are consistent with the pattern found in the subset of ports for which we do observe the total

¹²This follows Ducruet et al. (2018).

¹³Nightlight data suffer from the well-known “bleeding” problem; high levels of luminosity spill over into neighboring cells biasing the estimated area of a city systematically upward (Donaldson and Storeygard, 2016). To circumvent this problem, we use data on the area of the city available for a subset of our city sample from Geopolis to extrapolate the area of the city in a specification that includes the area of the city as measured by the nightlight data and country fixed effects. To probe robustness, we also estimated the area of the city based on daylight satellite data following Vogel et al. (2019), with similar results.

port area.

3 The reduced form effects of new port technologies

In this section, we estimate how the arrival of new port technologies in the 1960s affected the location of economic activity. We start by documenting a novel stylized fact; beginning in the 1960s, ports reallocated from the largest cities to smaller ones. We show that this effect was driven by initially smaller cities increasing their shipping activities disproportionately relative to large ones. We hypothesize that this reallocation is driven in large part by the fact that as new port technologies require more land, it became too expensive to host ports in the largest, highest land price cities. We provide several pieces of evidence in support of this claim.

3.1 Reallocation of ports from big to smaller cities

The qualitative literature suggests that one effect of new port technologies was the reallocation of ports from large to smaller cities (Levinson, 2010). Using our data, we are able to examine this relationship across our worldwide set of port cities. We begin by estimating the correlation between population and shipping separately by decade in Table 2 in a specification that includes year (columns 1 and 3) or country-by-year fixed effects (columns 2 and 4). There is a strong, positive correlation between shipping and population throughout our sample period. This is unsurprising as larger centers of economic activity tend to have higher demand for shipping. More interestingly, this correlation begins weakening in 1970, at exactly the same time as new port technologies are introduced according to the literature. The decline in correlation is statistically significant between 1960 and 1990 and the magnitude is large; the size of the coefficient decreases by 20% over this time period. Moreover, it does not seem to be driven by regional trends such as the rise of Asia in the late 20th century, as the pattern becomes stronger when only within country variation is used.

These results suggest that port services reallocated from the largest to smaller cities at the same time as new port technologies were adopted worldwide. One issue with this interpretation is that our baseline sample only includes shipping that takes place in ports located within the city (port cities). It is possible that the weakening correlation could be driven not by shipping reallocating to smaller cities that we observe in our sample, but by large cities moving their port to a small location that is close to, but not within

the city, and outside of our sample of cities.¹⁴ To address this, Table 11 re-estimates the specification adding shipping flows that serve a city from outside. In particular, columns (3)-(4) add additional shipping flows from standalone ports (not assigned to a city in our baseline sample) which serve the city from outside. The weakening correlation is apparent in these specifications and similar in magnitude. Based on this, we conclude that the weakening correlation is indeed driven by port services reallocating from the largest to smaller cities, rather than by the reallocation of shipping activities to a small location just outside the city.

How did this reallocation take place? Figure 4 splits the sample into size quartiles based on initial population in 1950 and estimates the growth in shipping over time. In the decades after 1960, there is a diverging trend in growth by size quartiles. The initially smallest three quartiles of cities increase their shipping flows the most, while there is only relatively modest growth in shipping flows of the initially largest quartile of cities. While convergence in shipping flows across ports is a possible explanation, trends only start to diverge in 1970, when new port technologies were introduced, making convergence an unlikely explanation.

While the timing of these trends suggests that new port technologies may be a driving force, we perform a stronger test based on exogenous variation in a city's suitability to new port technologies. In particular, we show that if a city is more suitable to new technologies for exogenous reasons, it experiences faster shipping growth if its initial population is small. To this end, we introduce a novel measure of exogenous suitability to new port technologies that we will use throughout the paper.

Exogenous port suitability

We use the area of deep sea around, but not directly at the city to isolate exogenous variation in port suitability. A large area of depth around the city became a locational advantage for modern ports for at least two reasons. First, ships have become vastly bigger, requiring greater depth (Brooks et al., 2018). Second, depth over a large area has become important as modern ports achieve part of their cost-efficiency through fast turnaround times for cargo, and through scale. Up to the 1950s, a ship would often spend weeks berthing in port, while today, ships turn around in as little as 24 hours (Hoffmann and Sirimanne, 2017; Levinson, 2010). This requires ships (many of which operate on regular schedules similar to buses) to wait in the vicinity of the port before their berthing time starts. Depth over a large area allows for more ships to be berthed in parallel to achieve scale, and for ships to

¹⁴An example is London, UK which was one of the largest ports until the 1950s and is today served predominantly by Felixstowe, over 100km outside of the city.

wait in the direct vicinity of the port achieving faster turnaround times. Figure 5 illustrates this point by contrasting the distance at which ships are anchored waiting to berth in Los Angeles, a large, successful containerized port, and Buenos Aires, a top 10 port in terms of shipping flows in 1951, but one that lost significant ground by the 1990s. Figure 6, plotting depths around the two ports, makes clear why ships need to anchor much further from the port in Buenos Aires; differently from Los Angeles, as the sea around Buenos Aires is not very deep, ships waiting to berth have to remain anchored in deep water at a distance of about 50 kms from the actual port. This will increase turnaround times as ships cannot wait close to the port itself. This implies that even holding both ship and port sizes constant, successful modern ports need depth over a larger body of water to remain competitive.

Following Brooks et al. (2018), we think of depth as an exogenous cost-shifter that makes it cheaper for a port to reach a desired depth by investing in costly dredging. The empirical challenge is that observed port depth is a combination of naturally endowed depth and depth attained by dredging. To isolate the exogenous component of this, we construct a novel measure of depth around the port using granular data on underwater elevation levels.¹⁵ We take buffer rings of various sizes around the geocode of the port and sum the number of sea cells deeper than a particular threshold within each buffer. This measure captures two sources of variation that affect a city’s suitability to become a modern port. First, it captures the availability of large bodies of water around the city, which is important given the larger size of modern ports. Second, it also depends on the depth of the sea around the port, which affects the costliness of becoming a modern deep-sea port.

To operationalize this measure of naturally endowed port depth, we need to choose the buffer and the threshold depth. As there is no guidance for picking these, our approach is to pick a baseline and examine the sensitivity of our estimates based on changing the buffer or threshold. Our baseline measure of exogenous port suitability (*port suitability*, henceforth) is 1 plus the log of the sum of cells in the 5-10 km buffer ring around the port that are deeper than 30 feet.¹⁶ We plot the spatial distribution of port suitability on Figure 7. Reassuringly, there is significant variation even for cities relatively close to each other. However, we test for spatial correlation in robustness checks.

¹⁵Data on underwater elevation levels are from the Gridded Bathymetric Chart of the Oceans.

¹⁶The depth measure used in Brooks et al. (2018) also uses a 30 foot cutoff, though depth in that paper is measured at the port in 1953 as reported in the World Port Index. Note that while Brooks et al. (2018) use depth at the port to predict which ports get containerized in the US (a binary indicator), this type of point-based depth measure does not predict shipping flows well in our worldwide dataset. This is possibly due to the fact that the area based measure used in this paper has better predictive power for whether the port is able to achieve the scale required for becoming a large modern port.

One concern with this source of variation is that natural port depth has always conferred some advantage to a location, as it lessened the requirement for costly dredging. While this is certainly true, until the advent of modern port technologies, depth was not required across such a large area, implying that it became far more costly to attain a particular depth for the entire port after the advent of modern port technologies.¹⁷

Given we observe shipping flows from 1950 onwards, we can test the extent to which our measure of port suitability explains shipping in the period after, but not before the arrival of new port technologies. We estimate the effect of port suitability on shipping flows allowing the coefficient to differ by decade in a specification that includes city and year fixed effects in column (1) of Table 3. Cities that have better port suitability witness faster growth in shipping in the decades following 1960, but not before. The coefficient on the interaction between port suitability and the 1960 decade dummy is small and not statistically significant. In the decades following 1960, once new port technologies were being adopted worldwide, cities with higher port suitability witnessed faster growth in shipping flows. The coefficients increase in size by more than an order of magnitude relative to 1960 and are highly statistically significant. Column (2) estimates the effect of port suitability on the growth of shipping flows after 1960 with a single coefficient that measures the effect of port suitability on shipping flows after containerization relative to before. Consistent with column (2), shipping increased disproportionately in cities with higher port suitability after 1960. Table 12 in the online appendix contains robustness checks to the buffer used and the variation included in the port suitability measure.¹⁸ In the remainder of the paper, we use our baseline measure of port suitability to isolate exogenous variation in shipping flows.

Returning to our specification of interest, we want to understand the extent to which differential shipping growth across initial population size quartiles has a causal interpretation. To answer this question, we examine heterogeneity in the effect of port suitability on shipping by size quartile. Column (3) re-estimates the effect of port suitability on the growth of shipping flows after 1960 (as in column (2)), but allowing for the effect to vary by initial size quartile. The effect on the initially smallest three quartiles of cities is the largest and highly significant. The three quartiles are similar in magnitude. However, there is a markedly smaller effect on the initially largest cities. The coefficient is almost halved in size and is

¹⁷For example, MacElwee (1925) writes that New York, one of the largest ports of the world at the time, had a dredged channel of 40 feet to only some, but by no means all parts of the harbor.

¹⁸Port suitability as measured by the sum of cells above 30 feet deep has similar predictive power at buffers between 3-15km. The measure does not have sufficient predictive power at closer buffers, and it's strength falls farther away from the port.

statistically different from the smallest quartile at 10% significance. This implies that cities with the same exogenous suitability to new port technologies respond differently based on their initial size.

Taking these results together, ports reallocated from large to smaller cities in the period after the arrival of new port technologies. This was accounted for by initially smaller cities increasing their shipping flows disproportionately relative to the initially largest ones. These effects are partly driven by the fact that cities equally well-suited to modern port technologies responded more to the arrival of new port technologies if they were smaller in size. Why would the arrival of new port technologies lead to this type of reallocation of ports? Informed by the nature of new port technologies, we hypothesize that this is due to a land-price mechanism. As new port technologies use land more intensively than the technology they replace, it became more costly to host port services in the largest, most congested cities where the opportunity cost of land was the greatest. The following section provides several pieces of evidence to support this claim.

3.2 Reallocation driven by the land mechanism

First, we examine the extent to which shipping responds differentially based on the elasticity of land supply in a city. If our hypothesis is correct, setting up port services should be more attractive in cities where land is cheaper and where it is easier for the city to expand. To test this mechanism, we proxy for the price of land by constructing the Saiz land supply measure for our sample of port cities (Saiz, 2010). This measure captures the extent to which a city is constrained by natural barriers to expanding in size such as the presence of bodies of water, or high elevation.¹⁹ We estimate heterogeneous effects in land supply elasticity by including an interaction term of port suitability and the Saiz land supply elasticity to the specification estimated in column (2) of Table 3 (which switches on only in the decades including and after 1970). We plot the marginal effect of port suitability on shipping evaluated at different values of the Saiz land supply elasticity measure in Figure 8. There is a considerable amount of heterogeneity in land supply elasticity. At low values of the land supply elasticity (around 0), the effect of port suitability on shipping is about 0.1 and marginally insignificant at 10%,

¹⁹More formally, we replicate the Saiz measure as closely as possible by summing the number of cells in a circle of radius 50km where the city can potentially expand. This circle excludes bodies of water and overland cells where the elevation slope is greater than 15%. This measure is naturally bounded between 0 and 1, with higher values corresponding to a larger share of cells along which the city can expand. In practice, the Saiz measure in our sample of cities spans almost the whole range of values between 0 and 1, with a median of 0.51.

which is roughly half the size of the baseline effect of 0.21 in column (2) of Table 3. At intermediate values (around 0.5), the effect is around 0.2, highly statistically significant and roughly the same size as the baseline coefficient. At the highest values of the Saiz measure (around 1), the effect is larger than 0.3, or about 50% larger than the baseline effect. Based on this, shipping increased significantly more in cities that were equally suitable for modern port activities based on the amount of deep sea area available in the vicinity of the city, but where land was in more elastic supply, consistent with our hypothesized mechanism.

Second, if the land price mechanism is at work after the arrival of new port technologies, the *share* of land allocated to the port should grow differentially between large and small cities. Larger, more congested cities tend to have higher land prices. Hence, all else equal, larger cities should allocate less land to ports according to our hypothesis. We investigate this by plotting the binscatter of the share of land allocated to the port in Figure 9 against its concurrent population level in the years 1950-1990. Contrary to the strong, positive correlation evident in the data between shipping and population throughout our sample period, we find that there is a negative correlation between the share of land allocated to the port and its population in 1990. That is, larger cities dedicate a smaller *share* of land to their port, despite the fact that they have larger shipping flows as discussed in the previous section. Moreover, this has not always been the case. The correlation between port area and population was strongly positive in the 1950s and 1960s and weakened thereafter before turning negative in the 1980s.²⁰ While these estimates do not have a causal interpretation, they do suggest that large cities are relatively less specialized in port services than smaller cities in terms of their land allocation.

Finally, we examine the local, reduced form effects of shipping. Improvements in transportation infrastructure can have positive local effects by making the location more attractive for firms and consumers. However, our land price mechanism suggests that hosting transportation infrastructure may also be costly for a location, as it crowds out other productive activities through the vast amounts of land that modern port activities require. Our empirical setting provides an indirect way of testing this. First, we estimate the local effect of shipping on economic activity in the port city using our panel of shipping flows for 1950-1990. Our proxy for economic activity is the population of the city. We estimate a standard panel specification of population in city i at time t on contemporaneous shipping flows where we include city and time fixed effects. We isolate exogenous variation in shipping using our port

²⁰Table 4 presents the regression coefficients for these specifications, both in the log of population and in levels.

suitability measure interacted with a binary indicator that takes the value of one in decades after and including 1970. Table 5 includes the baseline estimation results. Both the OLS and the 2SLS estimates are positive and highly significant, with the 2SLS estimates being somewhat larger. According to the 2SLS estimate, doubling shipping increases population by 14% on average. The first stage and reduced form estimates (column 3 and 4) are of the expected positive sign and highly statistically significant. The KP F-stat for the first stage is 28. Columns (5) and (6) show the pre-trends test for both the reduced form and the first stage; the effect of port suitability on shipping and population before the adoption of modern port technologies is small and statistically not different from zero. The coefficient becomes large, positive and highly significant in 1970 and stays high throughout the sample period.

In line with previous findings in the literature (Brooks et al., 2018), the local effect of shipping on economic activity is positive.²¹ Having established that our estimated effect on port cities is similar to findings in the previous literature, we turn to examining whether there is evidence for the land-price mechanism. We test for this in the following way; cities near a port benefit similarly (although not equivalently) from improvements in market access when ports get bigger, but would not face the same cost in terms of allocating scarce land to the port. If the cost effect is sufficiently strong, we would expect the estimated effect on nearby inland cities to be *larger* than that on port cities. We assign inland cities to their nearest port (reachable overland) and estimate the same specification as for port cities. We estimate the effect on distance intervals of 100 kilometres in rolling windows of 25 kilometers (the first specification includes all cities 0-100km from their nearest port, the next 25-125km from their nearest port and so forth). Figure 10 displays the 2SLS coefficient estimates and their 10% confidence intervals for the port city (own effect) and the effect on nearby inland cities. The results are striking. The estimated effect for nearby inland cities is consistently *larger* than that for the port city until the 100-200km buffer is reached, though the confidence intervals do overlap. At 100-200km, the coefficient starts shrinking towards zero and is no longer statistically significant.²²

²¹We subject these results to a variety of robustness checks. We drop continents one at a time (Figure 15), test robustness of the results to the inclusion of non-parametric time trends by continent, ocean and initial size groups (Table 13), and estimate the long difference (Table 14 and Figure 16).

²²Table 6 contains the estimates for the specifications plotted in Figure 10, including the pre-trends tests for each specification. It should be noted that these specifications include non parametric continent-by-year time trends. The reason for this is that inland cities are vastly over-represented for some continents as Figure 2 shows. For this reason, pre-trends are apparent in the specifications that do not include continent-by-year trends. Figure 14 and Table 15 show that while the result that the effect on nearby inland cities is larger than the port city holds for these specifications, these should not be interpreted as causal because of the

This section has presented three pieces of evidence that support our hypothesized land-price mechanism. While there are alternative mechanisms that may explain each separately, we view the three together as providing strong, reduced form evidence for the hypothesis that the land price mechanism is partly driving the reallocation of ports from big to smaller cities. In particular, the larger estimated effect on inland cities is causal evidence that suggests port cities face a cost of shipping activities that does not affect nearby inland cities in the same way. While this piece of evidence is silent on what that cost could be, the other two pieces of evidence speak directly to the land-price mechanism.

In summary, while the local (net) benefit of shipping are positive, there seem to be important costs associated with investing in modern port infrastructure. Based on the findings presented in this section, these costs are arguably driven by the fact that modern ports require vastly more land, which make port activities particularly costly in bigger, higher land price cities. Separating the costs and benefits of port infrastructure investments using only reduced form empirical methods is difficult however. For this reason, the next section presents a quantitative, general-equilibrium spatial model that incorporates both the benefits of investing in port infrastructure by way of improved market access, but also the costs it imposes on host cities through the land market. As we show below, the model suggests an empirical strategy for separating the costs and benefits. Additionally, it is sufficiently rich to allow us to quantitatively estimate the effects of new port technologies in general equilibrium.

4 A model of cities and endogenous port development

To study the costs and benefits of port development as well as the aggregate effects of port development induced by new port technologies, we develop a 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. Developing the port, however, is costly as it requires scarce land that can be used for other purposes. In equilibrium, cities will differ in the benefits and costs of developing the port, which will make certain cities gain relative to others and trigger a reallocation of both shipping and overall economic activity across cities.

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

violation of pre-trends.

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 inland 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

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

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.

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 et al. (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 in each city to one. Landlords have the same preferences 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 through the port 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 and hence profits from transshipment services are zero in equilibrium.²³

4.1.3 Production

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

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

where $q(r)$ denotes the firm's output, $\tilde{A}(r)$ is total factor productivity in the city, $n(r)$ is the amount of labor used 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 inland cities). Hence, $1 - F(r)$ is the remainder of land that landlords rent out to firms for production.

We incorporate agglomeration economies by assuming that total factor productivity depends 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 $\alpha \geq 0$ is a parameter that captures the strength of agglomeration economies. The representative firm does not internalize the effect that its labor demand decision has on local population. Hence, it takes $N(r)$ as given.

²³We assume that the stocks of land owned by the various landlords in a given port city do not differ in quality or proximity to water, hence they are equally suitable to transshipment services. However, we incorporate the fact that port cities are differentially suitable to transshipment services (for instance, they have different natural port depth) by incorporating an exogenous part of transshipment costs that varies across cities.

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:²⁴

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

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

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

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, \dots, p_M takes the form

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

where $\phi_\zeta(d(r, p_0))$ corresponds to the overland shipping cost between the origin and the first

²⁴The assumption of idiosyncratic shipping cost shifters follows Allen and Atkin (2017) 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 (2017), finding optimal shipping flows is computationally more demanding.

port en route p_0 , and $\phi_\zeta(d(p_M, s))$ 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 .²⁵

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.²⁶ 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 the transshipment services is driven down to marginal cost and therefore

$$O(p_m) = [\nu(p_m) + \psi(F(p_m))] Shipping(p_m)^\lambda$$

in equilibrium.

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. Finally, firms choose their production and shipping of goods to maximize their profits, taking prices,

²⁵Note 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.

²⁶To 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.

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.²⁷

4.2 Predictions of the model

This section presents the qualitative predictions that the equilibrium conditions of the model deliver on port development and its consequences on the spatial distribution of population across cities. These predictions follow from the equilibrium conditions of the model, which are derived in the Theory appendix.

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

$$-\psi'(F(r)) = \frac{R(r)}{\text{Shipping}(r)^{1+\lambda}} \quad (1)$$

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

$$R(r) = \frac{1 - \gamma w(r) N(r)}{\gamma (1 - F(r))} \quad (2)$$

such that $w(r)$ is the wage in city r . As the left-hand side of equation (1) is decreasing in $F(r)$ 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 port 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. The two effects together imply that new port technologies that lower shipping costs and increase shipping flows lead to more land allocated to the port, in line with what we have documented in the data.

²⁷This equilibrium definition implies that we do not give landlords the right to choose the amount of transshipment they conduct; in other words, landlords cannot refuse the provision of transshipment services to anyone at the market price. This assumption is needed for computational tractability, as it allows us to abstract from a corner solution in which the supply of transshipment services is zero. In line with this logic, we can relax the assumption and allow landlords to choose any *positive* amount of transshipment, but not zero transshipment. Generalizing the model this way does not change the equilibrium as landlords' profits are linear in the amount of transshipment and zero in equilibrium, hence landlords are indifferent between transshipping any two amounts as long as they are both positive.

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

Proposition 2 highlights that the cost of adopting new port technologies differs across cities. Those cities that have high land rents do not allocate much land to the port as the opportunity cost of land is very high. As cities with smaller population tend to have lower rents by equation (2), this finding suggests that port development must have happened primarily in small cities. Hence, new port technologies should have triggered a reallocation of shipping from large to small port cities, consistent with what we have documented in the data.

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

$$N(r)^{[1+\eta\sigma+(1-\gamma-\alpha)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} = \gamma^{\sigma-1}\tilde{a}(r)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} A(r)^{\frac{(\sigma-1)^2}{2\sigma-1}} (1-F(r))^{(1-\gamma)\frac{(\sigma-1)^2}{2\sigma-1}} MA(r) \quad (3)$$

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.

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: a *market access effect* that increases the population of the city, and a *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).*

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 (3) in particular, provide a structure that allows 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 re-estimate the effect of port development on economic activity guided by the insights of the model presented in the previous section. Section 3 showed that the reduced form effect of port development on population was positive in port cities. The model suggests two opposing forces at work; the direct effect of port development is to crowd out population as increased competition for scarce land drives up rents, while the indirect effect is a decrease in transportation costs that improves a city's market access, thereby drawing population in. We examine whether the data are consistent with the forces in the model by estimating the following specification:

$$POP_{it} = \alpha_i + \delta_t + \phi_1 * SHIP_{it} + \phi_2 * MA_{it} + \epsilon_{it} \quad (4)$$

where $MA_{it} = \ln \left(\sum_{s=1}^S \frac{POP_{st}^{[1-\eta(\sigma-1)-(1-\gamma-\alpha)\sigma]\frac{\sigma-1}{2\sigma-1}}}{\mathbf{E}[T_t(i,s)]^{\sigma-1}} \right)$ is the empirical equivalent of the model-based market access term.

The difference with respect to the reduced form estimating equation presented in Section 3 is that in light of the model, once we control for the market access of a city, ϕ_1 should recover the direct effect of increased shipping activity, which we expect to be negative. In contrast, we expect ϕ_2 , the elasticity of population to market access, to be positive. We view this specification as a test for the model forces as the shipping term captures any (positive or negative) effect of shipping on population above and beyond the effect stemming from improved market access.

It should be noted that estimating equation (4) does not exactly correspond to equation (3) in the model. The key difference is that we proxy time-varying port size using shipping flows to estimate the direct effect of port development (ϕ_1), and we do not include port size in the market access term. However, even the simplified version of the estimating equation presents a number of empirical challenges.

First, there is a question regarding what the baseline sample should be. The model

applies to all cities, regardless of whether they are inland or 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-s used to identify ϕ_1 and ϕ_2 (explained below) are only defined for port cities. Accounting for inland 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 city dataset, the market access of port cities is calculated using the full set of (inland and port) cities.

Second, time-varying bilateral trade costs $\mathbf{E}[T(i, s)]$ between origin and destination are not observed in the data. To overcome this challenge, we estimate time-varying bilateral iceberg trade costs between all cities (inland and port) in our dataset. More specifically, we incorporate three types of costs: 1. the cost of shipping overland, 2. the cost of sea shipping, and 3. the transshipment costs associated with crossing ports. We use the following parameters to calculate the lowest shipping costs between each pair of cities using a fast-marching algorithm. Based on Allen and Arkolakis (2014), we assume that the overland shipping cost ϕ_ζ and the sea shipping cost ϕ_τ take the form

$$\phi_\zeta(d) = e^{t_\zeta d} \quad \phi_\tau(d) = e^{t_\tau d}$$

where d is distance traveled, and set the values of t_ζ and t_τ to the corresponding road and sea shipping cost elasticities estimated by Allen and Arkolakis (2014).²⁸

There are no readily available measures of transshipment costs that we are aware of. To construct these, we use the following approach. Both the model and the literature on port development suggest that larger ports are more efficient. We use port efficiency measures, available for a subset of our sample from Blonigen and Wilson (2008), to estimate the empirical relationship between port efficiency and shipping flows at the port level in our data. We use the estimated coefficient from this regression to predict port efficiency for all the ports in our data across each decade.²⁹ Note that changing transshipment costs are the

²⁸Allen and Arkolakis (2014) also allow for fixed costs of inland and sea shipping. However, they set the fixed costs of road shipping to zero. In the case of sea shipping, our aim is to define transshipment costs in a broad sense such that they include any cost that is not a function of shipping distance, such as fixed costs of sea transportation.

²⁹The port efficiency measures are estimated as exporter port fixed effects in a regression of bilateral HS 6-digit product level import charges that control for distance, value, value-to-weight, percentage of containerized traffic between the two ports, trade imbalances, time, product and importer port fixed effects using US census data. The exporter fixed effects are all estimated relative to the efficiency of the port of Rotterdam. For our purposes, these relative port efficiencies need to be scaled to levels. We do this by

only source of time series variation in our estimated trade costs.

Third, the model-based measure of market access requires taking a stand on the values of the parameters η , σ , γ and α . Table 7 contains the parameter values we use and their source. As we use the same values when taking the full model to the data, Section 6.3 discusses the calibration of all structural parameters in detail.

Identifying ϕ_1 and ϕ_2 requires two sources of exogenous variation. The two IV-s we propose build on our measure of exogenous port suitability. We use the baseline measure of depth in the vicinity of the port as one instrument and construct a second instrument for market access based on the insight that the depth of other ports reduces bilateral trade costs and hence affects the market access term. In particular, other cities' port depth should affect bilateral trade costs from that city to all of its trading partners. Based on this reasoning, our second instrument is defined as follows:

$$MAIV_{it} = \sum_s \frac{POP_{s,1950}^{[1-\eta(\sigma-1)-(1-\gamma-\alpha)\sigma] \frac{\sigma-1}{2\sigma-1}}}{\left(\widehat{T}_t(i, s)\right)^{\sigma-1}}$$

where \widehat{T}_t , predicted transport cost, is predicted based on bilateral distance, and depth if year ≥ 1970 . That is, we replace the bilateral trade cost terms in the market access measure with the predicted decrease in transport costs between i and s driven by the exogenous port suitability of trading partner s .

Table 8 presents the estimation results. Columns (1) to (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 smaller 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. The instruments are strong, yielding a combined Kleibergen-Paap F-statistic of 14.27.³⁰

In summary, while the reduced form estimates show a net benefit of shipping on economic activity, based on the model inspired regressions, this is the result of two opposing forces. On

setting the iceberg trade cost of passing through Rotterdam to be 1.004. This is based on industry reports of port handling costs of \$140 AUD for a container and on a 20,000 EUR estimate of the average value of a container. In the fast marching algorithm, all ports in our sample, not only the ones that belong to a city with greater than 100,000 inhabitants are included. This is to ensure that we allow transshipment also through these smaller stand-alone ports, which may be transshipment hubs.

³⁰The online appendix shows that these results are robust to adding controls for non-parametric time trends in initial city size, continent or ocean (Table 16). Additionally, we show that the results are also robust to dropping cities in the close vicinity of the port city in the market access IV (Table 17).

the one hand, new port technologies improve a location's market access, drawing population in. On the other hand, there seems to be a strong negative direct effect of shipping on economic activity, consistent with crowding out. We conclude that this lends well-identified evidence for the model mechanisms. In the next section, we therefore turn to taking the full model to the data.

6 Taking the model to the data

Armed with evidence for the model's main mechanisms, we take the full structure of the model to the data in this section. This allows us to draw aggregate conclusions about the effects of changing port technologies in Section 8.

6.1 Backing out city-specific fundamentals

We combine data on city population, shipping flow and GDP per capita with the structure of the model to find the set of city amenities, productivities and exogenous transshipment costs $\nu(r)$ that rationalize the data. As the city population data are available for our 2,636 Geopolis cities, we choose these cities as the units of analysis. We observe shipping flows for the subset of Geopolis cities that are port cities. Following the time-invariant definition of port cities discussed in Section 2, a port city is defined as being a city with positive shipping flows in at least one of the decades between 1950 and 1990. We define inland cities as the remaining set of Geopolis cities. This gives us 553 port cities and 2,083 inland cities.³¹

City-level GDP per capita data are not readily available for our set of cities. We estimate GDP per capita for our full sample of 2,636 worldwide cities in the following way. We purchased 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 1990.³² We predict GDP per capita for our full sample using the linear fit of GDP per capita data from Canback 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).³³

³¹See Section 2 for further details on the city population and shipping data.

³²In particular, we merge the Canback data with our city list, and construct GDP per capita from the level of GDP and the population data provided by Canback. GDP are reported at purchasing power parity (in 2005 USD).

³³Most of the papers in this literature estimate the level of GDP within a country, where the level of development is not as widely dispersed as across the world. To account for these differences and the way in

Given that our city level GDP data are only available for 1990, we choose to invert the model based on the 1990 distribution of population, shipping and GDP. Since this data is measured after the advent of new port technologies, our counterfactual will increase transshipment costs in a way that mimics the lack of these technologies. Hence, the effect of new port technologies can be assessed by comparing the counterfactual equilibrium (old technologies) to the 1990 equilibrium (new technologies).

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 the value of V is chosen to match the ratio of shipping to world GDP. The rationale behind choosing this 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 $V = 364$.³⁴

GDP per capita maps into wages as

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

according to the model, where structural parameter γ is calibrated to 0.84, as explained in Section 6.3.

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 the Theory appendix.³⁵

which they affect luminosity, we include country fixed effects.

³⁴As not all our port cities have a positive number of ships in 1990 but the model cannot rationalize zero shipping flows under finite values of city-specific fundamentals, we change $Ship(r)$ from zero to one in these cities.

³⁵The 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 set of city-specific fundamentals that rationalize the data may be unique.

6.2 Shipping costs

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

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

and setting the elasticities t_{ς} and t_{τ} to the corresponding estimates in Allen and Arkolakis (2014).³⁶

We also need to choose endogenous transshipment costs as a function of the share of land allocated to the port (*port share*, henceforth), $\psi(F)$. 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 (1) and (2). One simple function that satisfies all these restrictions is

$$\psi'(F) = 1 - F^{-\beta} \tag{5}$$

where $\beta > 0$ to guarantee $\psi' < 0$. We can obtain ψ by integrating equation (5) as

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

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

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 cannot be identified separately from them as a result. Thus, we set κ to its theoretical lower bound $\bar{\kappa}$ without loss of generality when taking the model to the data. However, we do change κ in our counterfactual when our aim is to change transshipment costs back to their level prior to the advent of new port technologies.

Given the role that β plays in driving the relationship between shipping and the port

³⁶See Section 5 for details.

share according to equations (1) and (5), we calibrate this parameter to match the correlation between these two variables in the data. This gives us a value of $\beta = 0.023$.

6.3 Structural parameters

We are left with choosing the values of the model’s six structural parameters. On the production side, we take the estimate of the strength of agglomeration externalities, $\alpha = 0.06$, from Ciccone and Hall (1996), which 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%. 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. 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 almost exactly corresponds to our period of investigation. Together with $\alpha = 0.06$, this suggests setting $\gamma = 0.84$.³⁷

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 standard estimates of the trade elasticity (Simonovska and Waugh, 2014).

Finally, we have two structural parameters influencing shipping costs. One is the dispersion of idiosyncratic shipping costs, which – together with the formulation of these costs – we take from Allen and Arkolakis (2019), setting $\theta = 203$. The other is the elasticity of transshipment costs with respect 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 7 summarizes the calibration of our structural parameters.

7 Model fit

This section evaluates our model’s quantitative performance on data that was not targeted in the calibration. In particular, we can assess the fit of the model to data on the share of land that cities allocate to the port, which is untargeted. Producing a good fit to the distribution

³⁷Another 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, strictly speaking, absent from our model but might matter for the quantitative results.

of port shares is essential as they are key drivers of our novel mechanism; the crowding-out effect. Hence, Table 9 provides an extensive comparison between the model-predicted distribution of port shares $F(r)$ and the port shares observed in the data.

As Panel A of Table 9 demonstrates, the correlation between predicted and actual port shares is high, irrespectively of whether we calculate it in levels, in logs, or as a rank correlation. To understand whether these high correlations are primarily driven by city population and GDP (which we match between the model and the data), we compute each city’s residual port share from a regression in which we include population and GDP on the right-hand side both in the model and in the data. The correlation between these residual port shares is even higher than the raw correlation, suggesting that the high correlation is not driven by the mechanical matching of population and GDP between the model and the data. Rather, it seems that the model is able to predict well the cities that serve as transshipment ports for their hinterland.

Next, we look at how predicted and actual port shares relate to other city-specific observables. Panel B of Table 9 shows the results. Note that we have chosen our structural parameter β to match the correlation between port share and shipping, so this moment is targeted in the calibration. Nonetheless, the model also does well in replicating the correlation of port share with city population and GDP, which are untargeted. This is important, as Section 3 argued that the negative correlation between port share and population in the current data provides evidence for the presence of the crowding-out mechanism driven by high land prices in large and expensive cities.

Finally, we investigate how concentrated the distributions of predicted and actual port shares are by plotting the Lorenz curves of these two distributions (Figure 11). The two curves are quite close, although the model provides a slightly too low concentration among large ports relative to the data.³⁸ In other words, the model stops short at fully replicating the rise of mega-hubs after the advent of new port technologies. This is, however, not surprising given the functional form of sea shipping costs. As sea shipping costs are exponential in distance, the triangle inequality holds for them, therefore shippers do not have an incentive to ship cargo through more than two ports between the origin and the destination. In other words, hubs that specialize in *sea transshipment* – transshipping cargo from ship to ship – do not arise in the model (apart from a few shippers getting good idiosyncratic draws for routes that involve more than two ports). However, anecdotal evidence suggests that many

³⁸This difference also shows up if we calculate the Gini coefficient to characterize the overall concentration of port shares: it is 0.736 in the data but lower, 0.625, in the model.

of the largest ports specialize in sea transshipment in the data. Hence, it is not surprising that the model, lacking these types of ports, implies too little concentration at the upper tail of the port share distribution.

8 The effects of changing port technologies

This section illustrates how our calibrated model can be used to assess the aggregate and distributional effects of changing transshipment technologies in ports. To this end, we consider a counterfactual in which we increase the intercept parameter of the endogenous transshipment cost function, κ ; and decrease the shape parameter of the endogenous transshipment cost function, β . Increasing the intercept parameter implies a uniform increase in transshipment costs, corresponding to the lower overall efficiency of old transshipment technologies.³⁹ Decreasing the shape parameter β implies that the extensive use of land can lower transshipment costs to a lesser extent, corresponding to the lower land intensity of old transshipment technologies.

As our aim for now is to illustrate the effects of changing transshipment technologies rather than to estimate the full effect of new port technologies, we change the two parameters of the endogenous transshipment function in a very simple way. In particular, we increase the intercept parameter by 10 percentage points and decrease the shape parameter by 50 percentage points. This leads to a realistic increase in aggregate shipping by 92.6% from the counterfactual to the baseline.⁴⁰ The Theory appendix describes the details of how we perform the counterfactual simulation.

Table 10 summarizes the aggregate results of our counterfactual exercise. The last column of the table shows counterfactual values of variables subtracted from their baseline values. These should be compared to changes after the advent of new technologies in the data. Our counterfactual features a higher correlation between population and shipping (0.549) than the baseline (0.472), similar to the declining correlation we document in the data (Section 3). In other words, our counterfactual exercise can replicate the motivating empirical fact documented in this paper: the reallocation of shipping from larger to smaller cities. Moreover,

³⁹Recall that we take our model to 1990 data. Hence, our baseline world is one in which the new transshipment technologies are present, and we want to perform counterfactuals that mimic the lack of these new technologies.

⁴⁰As a reference point, the value of merchandise trade increased sixfold between 1950 and 1990 at constant prices. We have experimented with both smaller and larger changes in these parameters, and the effects seem monotonic. This suggests that there is nothing qualitatively special about the counterfactual we chose as our benchmark.

the relationship between port share and population becomes flatter in the counterfactual (a regression coefficient of -0.0052) than in the baseline (a coefficient of -0.0117), consistent with large cities becoming less specialized in port activities after the arrival of new port technologies.⁴¹

Aggregate world welfare increases by 5.3% from the counterfactual to the baseline, while aggregate real GDP increases by 5.7%. Estimates of the real GDP gains from trade tend to be substantially below these numbers in one-sector quantitative trade models. For example, Costinot and Rodríguez-Clare (2014) estimate a 2.28% effect of a 40% worldwide tariff (which, under their calibrated value of the trade elasticity, should correspond to a 200% change in trade), as well as a 4.4% effect of eliminating trade completely. These are much more dramatic changes in trade than the 92.6% increase in our counterfactual. A direct comparison of these numbers to our real GDP gains is not straightforward since standard quantitative models do not feature within-country geography. Keeping this caveat in mind, one interpretation of these results is that the larger effects that we find are due to the reallocation of economic activity caused by the arrival of new port technologies. Endogenous port infrastructure development has the potential to alter the size of the gains from trade if the costs are heterogeneous across locations, as is the case in our setting. This is due to the fact that as ports reallocate from large, high rent, productive cities, to smaller, lower rent and less productive cities, there is increased specialization based on comparative advantage. Larger cities specialize to a greater extent in their comparative advantage sector, non-port activities. Of course, if the spatial heterogeneity in costs plays a significant role in our welfare effects, we should find not only high aggregate gains but also large heterogeneity in the net gains across locations.

Hence, we look at heterogeneity in welfare gains implied by the model. As mobility is possible within countries, welfare equalizes within each country both in the baseline and in the counterfactual. Yet, welfare gains vary across countries. Figure 12 presents the distribution of these gains. As is apparent from the figure, the welfare effects of changes in transshipment technology vary dramatically across countries. Even though the majority of countries gain below 10%, some gain more than 40%. This heterogeneity substantially exceeds the heterogeneity across countries found by Costinot and Rodríguez-Clare (2014), who estimate that the welfare effects of eliminating trade completely vary between 1.5% and 8.1% in a

⁴¹Recall that the relationship is negative both in the baseline model and in the data. Contrary to what we find in the data, the relationship between port size and population is not positive in the counterfactual. This is unsurprising, as the counterfactual simulation still features a crowding out effect that is especially strong in large cities, albeit to a smaller extent than in the baseline.

standard one-sector quantitative model. The large cross-country heterogeneity we find lends additional support to the claim that our estimated aggregate benefits include efficiency gains that were induced by a reallocation of economic activity, which was triggered by heterogeneous local costs of infrastructure development. Hence, accounting for endogenous costs in port development seems to be an important margin for understanding the gains from trade and how they are distributed spatially.

9 Conclusion

TBA

References

- Abe, K. and J. S. Wilson (2009). *Weathering the storm: investing in port infrastructure to lower trade costs in East Asia*. The World Bank.
- Allcott, H. and D. Keniston (2017). Dutch disease or agglomeration? the local economic effects of natural resource booms in modern america. *The Review of Economic Studies* 85(2), 695–731.
- Allen, T. and C. Arkolakis (2014). Trade and the topography of the spatial economy. *The 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 (2017). Volatility and the gains from trade.
- Anderson, J. E. (1979). A theoretical foundation for the gravity equation. *The American Economic Review* 69(1), 106–116.
- Armenter, R., K. Miklós, and D. K. Nagy (2014). Bridges.
- Bernhofen, D. M., 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. A. and W. W. Wilson (2008). Port efficiency and trade flows. *Review of International Economics* 16(1), 21–36.
- Brancaccio, G., M. Kalouptsi, and T. Papageorgiou (2017). Geography, search frictions and endogenous trade costs. *National Bureau of Economic Research* (Working Paper 23581).
- Brooks, L., N. Gendron-Carrier, and G. Rua (2018). The local impact of containerization.
- Ciccone, A. and R. E. Hall (1993). Productivity and the density of economic activity. *NBER Working Paper* (No. w4313).
- Coşar, A. K. and P. Fajgelbaum (2016). Internal geography, international trade, and regional specialization. *American Economic Journal: Microeconomics* 1(8), 24–56.

- Combes, P.-P., G. Duranton, and L. Gobillon (2018). The costs of agglomeration: House and land prices in French cities. *Review of Economic Studies*, forthcoming.
- Corden, W. M. and J. P. Neary (1982). Booming sector and de-industrialisation in a small open economy. *The Economic Journal* 92(368), 825–848.
- Coşar, A. K., B. D. S. inside the box: Containerization, and trade (2018). Shipping inside the box: Containerization and trade. *Journal of International Economics* 114, 331–345.
- Costinot, A. and A. Rodríguez-Clare (2014). *Trade theory with numbers: Quantifying the consequences of globalization*, Volume 4 of *Handbook of international economics.*, pp. 197–261. Elsevier.
- 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. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review* 108(4-5), 899–934.
- Donaldson, D. and R. Hornbeck (2016). Railroads and american economic growth: A “market access” approach. *The 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.
- 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. and E. Schaal (2017). Optimal transport networks in spatial equilibrium.
- Gomtsyan, D. (2016). Rise of the machines: Evidence from the container revolution.
- Head, K. and T. Mayer (2014). *Gravity equations: Workhorse, toolkit, and cookbook*, Volume 4 of *Handbook of international economics*. Elsevier.
- Herrera Dappe, M. and A. Suárez-Alemán (2016). Competitiveness of south asia’s container ports: A comprehensive assessment of performance, drivers, and costs. directions in development. Technical report, World Bank.
- Hoffmann, J. and S. N. Sirimanne (2017). Review of maritime transport. *United Nations Publication*.
- Holmes, T. J. 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.
- Kennan, J. and J. R. 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.
- MacElwee, R. S. (1925). *Port Development*. McGraw-Hill.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6).
- Monte, F., S. J. 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. K. (2017). City location and economic development.
- OECD (2014). The competitiveness of global port cities. Technical report, OECD Publishing.
- Redding, S. J. and M. A. Turner (2015). *Transportation costs and the spatial organization of economic activity*, Volume 5 of *Handbook of regional and urban economics*, pp. 1339–1398. Elsevier.
- Rua, G. (2014). Diffusion of containerization.
- Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics* 125(3), 1253–1296.
- Santamaría, M. (2018). The gains from reshaping infrastructure: Evidence from the german division.
- Simonovska, I. and M. E. Waugh (2014). The elasticity of trade: Estimates and evidence. *Journal of international Economics* 92(1).
- Vogel, K. B., R. Goldblatt, G. H. Hanson, and A. K. Khandelwal (2019). Detecting urban markets with satellite imagery: An application to India. *Journal of Urban Economics*, forthcoming.

A Tables

Table 1: Comparison of sample of port cities and sample of cities with population (Geopolis)

	No port	Port	Total
No population information	0	1,592	1,592
Population information	2,083	553	2,636
Total	2,083	2,145	4,228

Table 2: Weakening correlation between population and shipping from 1960s

depvar	(1) ln(shipment)	(2) ln(shipment)	(3) ln(shipment)	(4) ln(shipment)
ln(pop)*1950	0.997*** (0.087)	1.330*** (0.098)	-0.064 (0.068)	-0.057 (0.076)
ln(pop)*1960	1.061*** (0.080)	1.386*** (0.095)		
ln(population)			1.061*** (0.080)	1.386*** (0.095)
ln(pop)*1970	1.022*** (0.080)	1.207*** (0.095)	-0.039 (0.068)	-0.179** (0.079)
ln(pop)*1980	0.981*** (0.078)	1.187*** (0.093)	-0.080 (0.079)	-0.199** (0.085)
ln(pop)*1990	0.860*** (0.070)	1.015*** (0.086)	-0.200** (0.087)	-0.372*** (0.094)
Observations	2,610	2,610	2,610	2,610
R-squared	0.203	0.541	0.203	0.541
Fixed Effect	year	country X year	year	country X year

Notes: Dependent variable: log of shipping flows at the level of the city. Regressor: log of population at the level of the city interacted by a decade dummy as indicated. Sample: port cities. Standard errors clustered at the city level. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Causal effect of new port technologies on shipping

VARIABLES	(1) ln(shipment)	(2) ln(shipment)	(3) ln(shipment)
Depth Measure X 1960	-0.010 (0.045)		
Depth Measure X 1970	0.225*** (0.049)		
Depth Measure X 1980	0.157*** (0.058)		
Depth Measure X 1990	0.236*** (0.065)		
Depth Measure X (≥ 1970)		0.211*** (0.041)	0.245*** (0.055)
Depth Measure X (≥ 1970) X Quartile 2			-0.049 (0.064)
Depth Measure X (≥ 1970) X Quartile 3			0.004 (0.069)
Depth Measure X (≥ 1970) X Quartile 4			-0.117* (0.061)
Observations	2,640	2,640	2,640
R-squared	0.130	0.129	0.132
Number of cityid	528	528	528

Notes: Dependent variable: log of shipping flows at the level of the city. Regressor: port suitability measured at the level of the city interacted by a decade dummy as indicated. “Quartile” indicates an interaction with a binary variable that takes the value of 1 if the city’s population in 1950 belongs to the specified quartile. Excluded category is quartile 1 (smallest). Sample: port cities. Standard errors clustered at the city level. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Correlation between share of a city occupied by the port and its population, over time

DEP VAR: port area/city area	(1)	(2)	(3)	(4)
population		0.0074** (0.0031)		
population * (year = 1950)	0.0074** (0.0031)			
population * (year = 1960)	0.0035* (0.0019)	-0.0039* (0.0023)		
population * (year = 1970)	0.0003 (0.0004)	-0.0070** (0.0028)		
population * (year = 1980)	-0.0002 (0.0001)	-0.0075** (0.0031)		
population * (year = 1990)	-0.0003** (0.0001)	-0.0076** (0.0031)		
ln(population)				0.0050*** (0.0016)
ln(population) * (year = 1950)			0.0050*** (0.0016)	
ln(population) * (year = 1960)			0.0041*** (0.0014)	-0.0008 (0.0006)
ln(population) * (year = 1970)			0.0026** (0.0010)	-0.0024** (0.0010)
ln(population) * (year = 1980)			0.0007 (0.0007)	-0.0043*** (0.0015)
ln(population) * (year = 1990)			-0.0003 (0.0007)	-0.0053*** (0.0016)
Observations	1,594	1,594	1,594	1,594
R-squared	0.0218	0.0218	0.0338	0.0338
Year FE	yes	yes	yes	yes

Notes: Dependent variable: share of land occupied by the port. Regressor: log of population / population at the level of the city interacted by a decade dummy as indicated. Sample: port cities where port area measures are available. Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Causal effect of shipping activity on population

	(1)	(2)	(3)	(4)	(5)	(6)
depvar	ln(pop)	ln(pop)	ln(shipment)	ln(pop)	ln(shipment)	ln(pop)
ln(shipment)	0.020*** (0.007)	0.140*** (0.051)				
Depth Measure X (\geq 1970)			0.217*** (0.041)	0.030*** (0.010)		
Depth Measure X 1960					-0.006 (0.046)	0.007 (0.006)
Depth Measure X 1970					0.233*** (0.049)	0.027*** (0.009)
Depth Measure X 1980					0.164*** (0.059)	0.032** (0.013)
Depth Measure X 1990					0.246*** (0.065)	0.042** (0.016)
Observations	2,609	2,609	2,609	2,609	2,609	2,609
R-squared	0.627	0.505	0.123	0.626	0.125	0.627
Number of cities	527	527	527	527	527	527
year FE	yes	yes	yes	yes	yes	yes
port FE	yes	yes	yes	yes	yes	yes
Type	OLS	2SLS	FS	RF	FS dynamic	RF dynamic
KP F-stat		27.85				

Notes: Dependent variable: log of shipping flows / population at the level of the city, as indicated. Regressor: log of shipping flows at the level of the city, “Depth Measure” indicates the port suitability measure interacted with decade dummy or indicator for decades including and after 1970, as indicated. Standard errors clustered at the city level. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Estimated causal effect of shipping on population of port city and nearby cities controlling for continent by year non-parametric time trends

Panel A: 2SLS estimates							
depvar	(1) ln(pop)	(2) ln(pop)	(3) ln(pop)	(4) ln(pop)	(5) ln(pop)	(6) ln(pop)	(7) ln(pop)
ln(shipment)	0.062 (0.039)	0.086** (0.036)	0.106** (0.042)	0.140** (0.057)	0.126** (0.059)	0.057 (0.068)	0.023 (0.058)
Observations	2,609	2,704	2,934	2,580	2,254	1,851	1,624
R-squared	0.731	0.633	0.591	0.522	0.593	0.742	0.794
Number of cities	527	547	596	524	460	381	336
sample	port city	0 - 100	25 - 125	50 - 150	75 - 175	100 - 200	125 - 225
KP	24.46	28.36	22.65	15.37	15.98	8.670	12.69

Panel B: Dynamic first stage							
depvar	(1) ln(shipment)	(2) ln(shipment)	(3) ln(shipment)	(4) ln(shipment)	(5) ln(shipment)	(6) ln(shipment)	(7) ln(shipment)
Depth Measure X 1960	-0.010 (0.048)	-0.062 (0.058)	-0.090 (0.057)	-0.073 (0.059)	-0.019 (0.059)	-0.085 (0.059)	-0.104* (0.062)
Depth Measure X 1970	0.224*** (0.052)	0.248*** (0.068)	0.233*** (0.064)	0.221*** (0.067)	0.216*** (0.064)	0.088 (0.068)	0.086 (0.070)
Depth Measure X 1980	0.142** (0.058)	0.184** (0.076)	0.103 (0.071)	0.087 (0.075)	0.107 (0.070)	0.073 (0.073)	0.149** (0.076)
Depth Measure X 1990	0.221*** (0.062)	0.392*** (0.081)	0.289*** (0.075)	0.237*** (0.077)	0.254*** (0.070)	0.182** (0.073)	0.203*** (0.076)
Observations	2,609	2,704	2,934	2,580	2,254	1,851	1,624
R-squared	0.168	0.177	0.192	0.194	0.204	0.221	0.232
Number of cities	527	547	596	524	460	381	336
sample	port city	0 - 100	25 - 125	50 - 150	75 - 175	100 - 200	125 - 225

Panel C: Dynamic reduced form							
depvar	(1) ln(pop)	(2) ln(pop)	(3) ln(pop)	(4) ln(pop)	(5) ln(pop)	(6) ln(pop)	(7) ln(pop)
Depth Measure X 1960	0.006 (0.005)	0.010 (0.007)	0.009 (0.006)	0.010 (0.007)	0.006 (0.007)	0.005 (0.008)	0.006 (0.008)
Depth Measure X 1970	0.019** (0.008)	0.030*** (0.010)	0.031*** (0.009)	0.035*** (0.010)	0.034*** (0.011)	0.021* (0.012)	0.020 (0.013)
Depth Measure X 1980	0.013 (0.010)	0.028** (0.012)	0.029** (0.012)	0.031** (0.012)	0.019 (0.013)	-0.002 (0.014)	-0.006 (0.015)
Depth Measure X 1990	0.014 (0.012)	0.035** (0.015)	0.035** (0.014)	0.041*** (0.014)	0.033** (0.015)	0.015 (0.016)	0.008 (0.017)
Observations	2,609	2,704	2,934	2,580	2,254	1,851	1,624
R-squared	0.749	0.744	0.743	0.766	0.763	0.779	0.801
Number of cities	527	547	596	524	460	381	336
sample	port city	0 - 100	25 - 125	50 - 150	75 - 175	100 - 200	125 - 225

Notes: Dependent variable: log of shipping flows / population at the level of the city, as indicated. Regressor: log of shipping flows at the level of the city, "Depth Measure" indicates the port suitability measure interacted with decade dummy, as indicated. Sample: port city indicates the baseline port city sample, 0-100 indicates the sample of cities 0-100km away from the nearest port (excludes port cities), 25-125 indicates the sample of cities 25-125 away from the nearest port and so forth. Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 7: 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)
$\sigma = 4$	Trade elasticity (Simonovska and Waugh, 2014)
$\theta = 203$	Idiosyncratic shipping cost dispersion (Allen and Arkolakis, 2019)
$\lambda = 0.074$	Congestion externalities in ports (Abe and Wilson, 2009)

Table 8: Model-based specification

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
depvar	ln(pop)	ln(pop)	ln(pop)	ln(pop)	ln(shipment)	ln(market access)	ln(pop)
ln(shipment)	0.023*** (0.007)	0.141*** (0.051)	0.006 (0.008)	-0.251*** (0.070)			
ln(market access)			1.670*** (0.632)	17.797*** (2.640)			
Depth Measure X 1970					0.208*** (0.041)	0.005*** (0.001)	0.038*** (0.009)
ivMA					-38.972*** (10.652)	1.372*** (0.301)	34.214*** (2.443)
Observations	2,571	2,571	2,571	2,571	2,571	2,571	2,571
R-squared	0.630	0.515	0.651	-1.346	0.133	0.702	0.683
Number of cityid	519	519	519	519	519	519	519
year FE	yes	yes	yes	yes	yes	yes	yes
port FE	yes	yes	yes	yes	yes	yes	yes
Type	OLS	2SLS	OLS	2SLS	FS	FS	RF
KP F-stat		27.96		14.27			

Notes: Dependent variable: log of shipping flows / population / market access at the level of the city, as indicated. Regressor: log of shipping flows at the level of the city, log of market access, “Depth Measure” indicates the port suitability measure interacted with a binary indicator for decades including and after 1970, “ivMA” is the IV for market access defined in the text. Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 9: Model fit

Panel A: Relationship between port size in the model and the data

Correlation (levels)	0.4218*
Correlation (logs)	0.5468*
Rank correlation	0.6277*
Correlation (residuals)	0.5801*

Panel B: Correlation between log port size and city-level observables

	Model	Data
Log shipping (targeted moment)	0.6060	0.6024
Log population	-0.1445	-0.0541
Log GDP	-0.2511	-0.0001

Table 10: Counterfactual results

Variable	Counterfactual	Baseline	Difference
Aggregate shipping	9,527,571	18,345,820	92.6%
Correlation of log population and log shipping	0.549	0.472	-0.077
Aggregate world welfare	21,165	22,281	5.3%
Aggregate world real GDP	429,018	453,682	5.7%
Coefficient of population on port share (millions)	-0.0052	-0.0117	-0.0065

B Figures

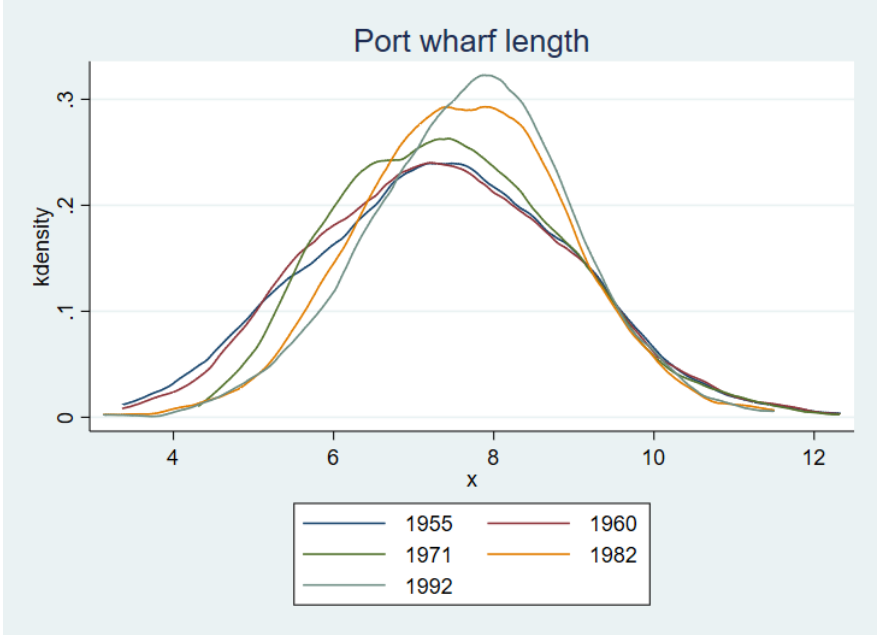


Figure 1: Port wharf lengths over time, across ports. Source: Ports of the World, 1955-1990

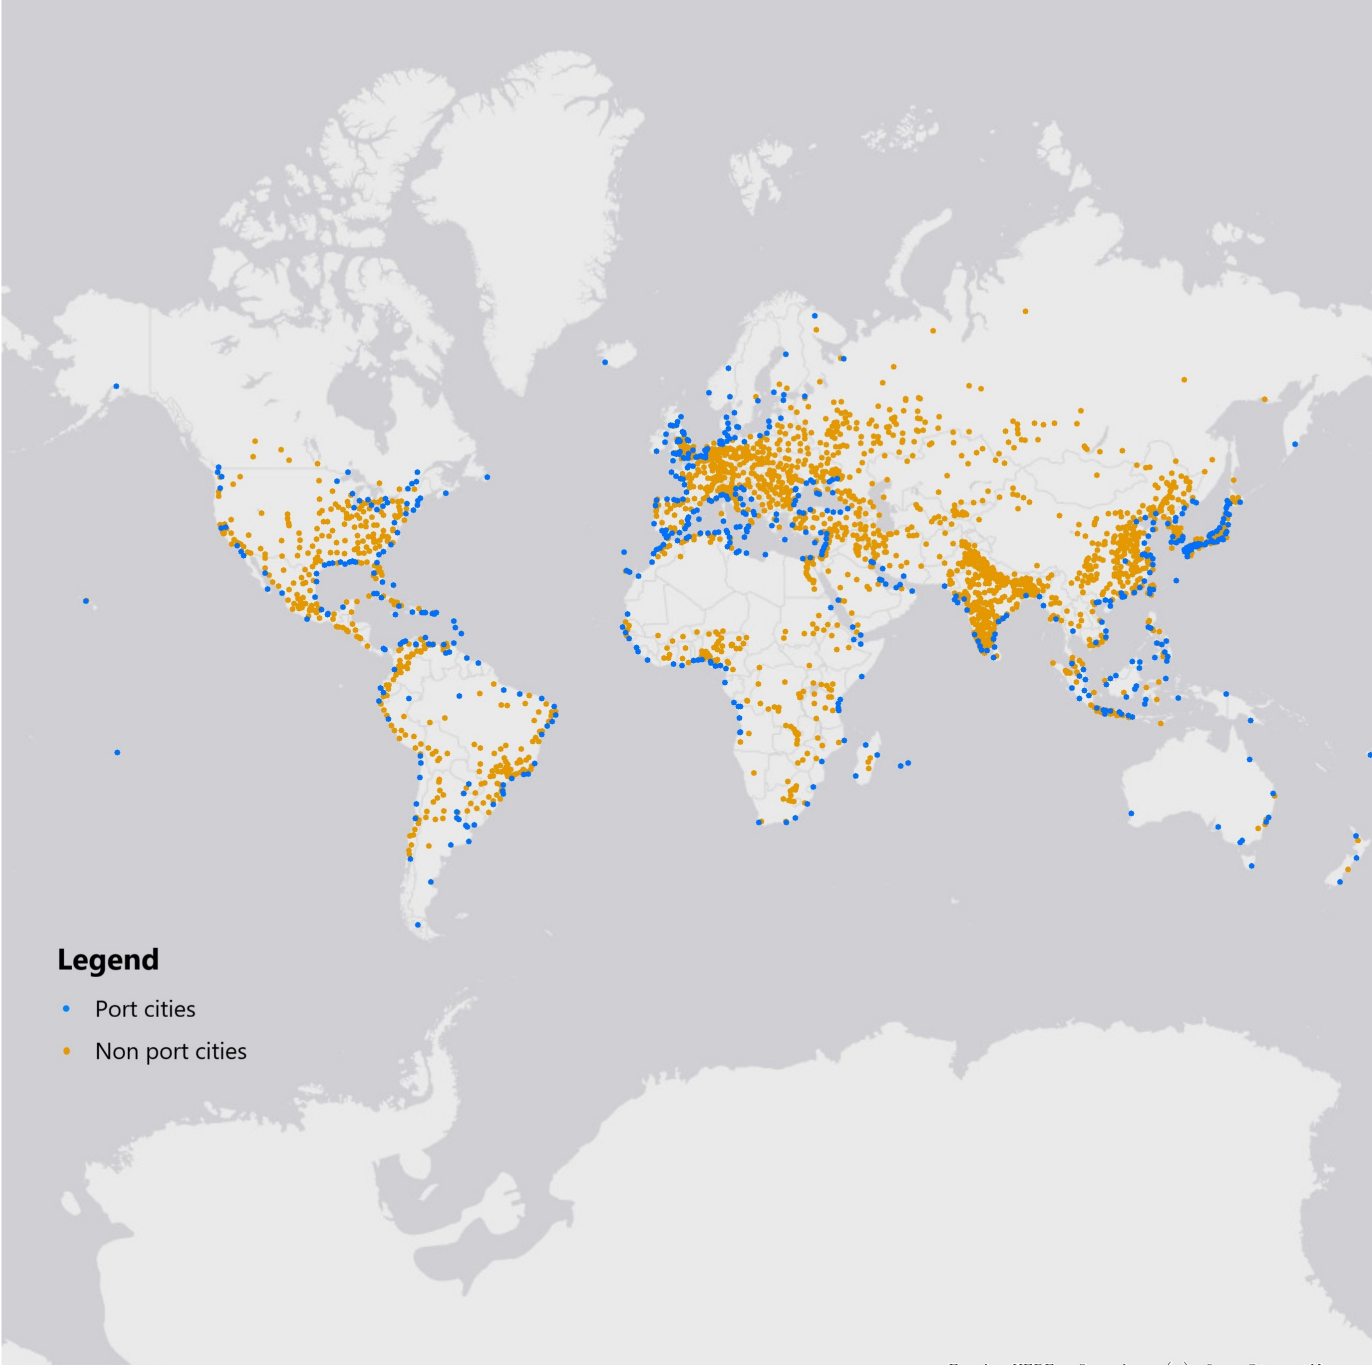


Figure 2: Port and non-port cities from Geopolis

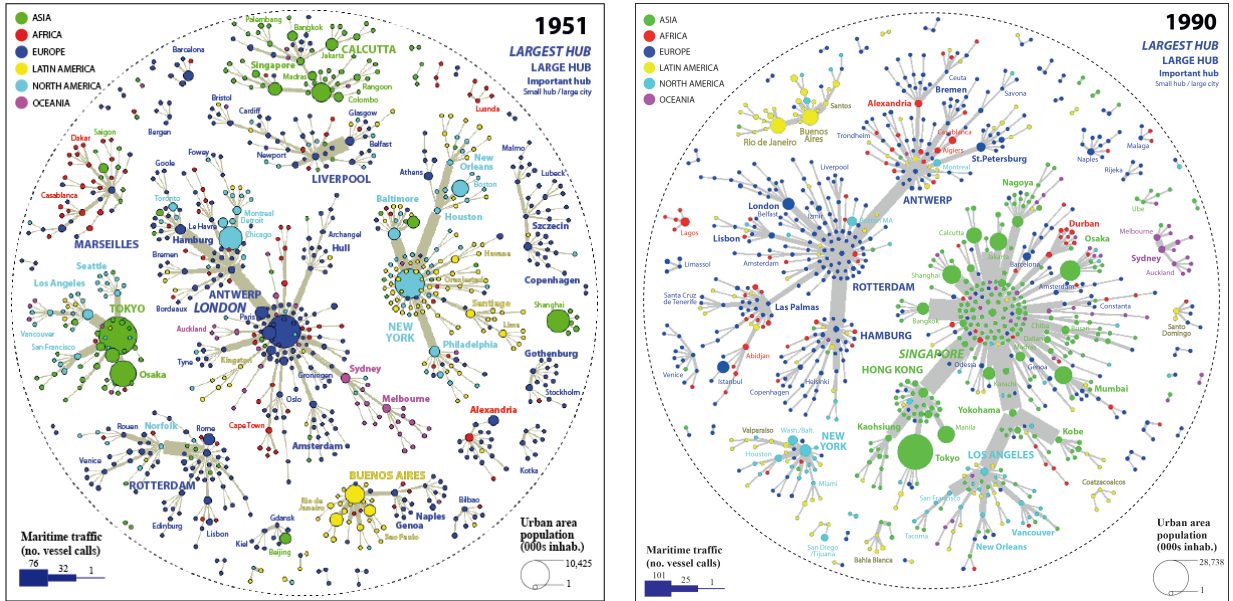


Figure 3: Shipping network in 1951 and 1990

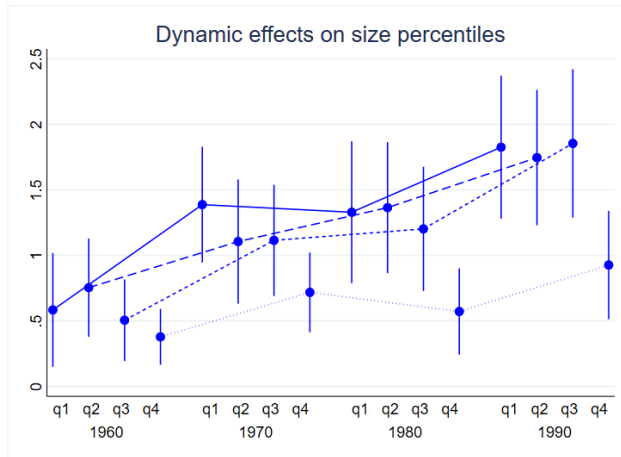
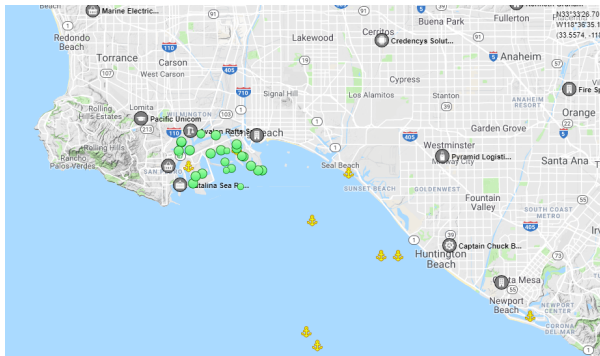
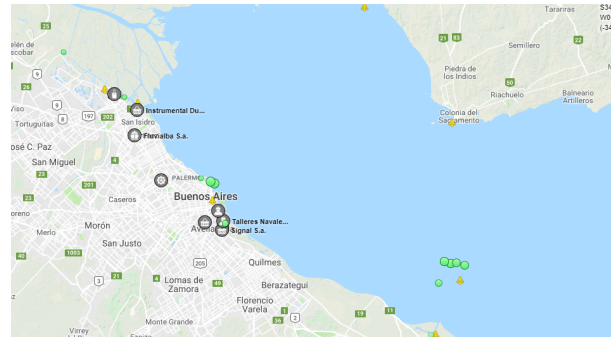


Figure 4: Dynamic effects on shipping, by initial population quartile

Notes: Estimated coefficients from a regression of the log of shipping flows at the city level on decade dummies interacted with the full set of initial size quartiles. Size quartiles defined based on population in 1950.




(a) Los Angeles




(b) Buenos Aires

Figure 5: Anchored cargo vessels, March 2019. Source: marinetraffic.com.
 Note: Figures not on same scale



(a) Los Angeles



(b) Buenos Aires

Figure 6: Elevation data (meters). Source: GEBCO

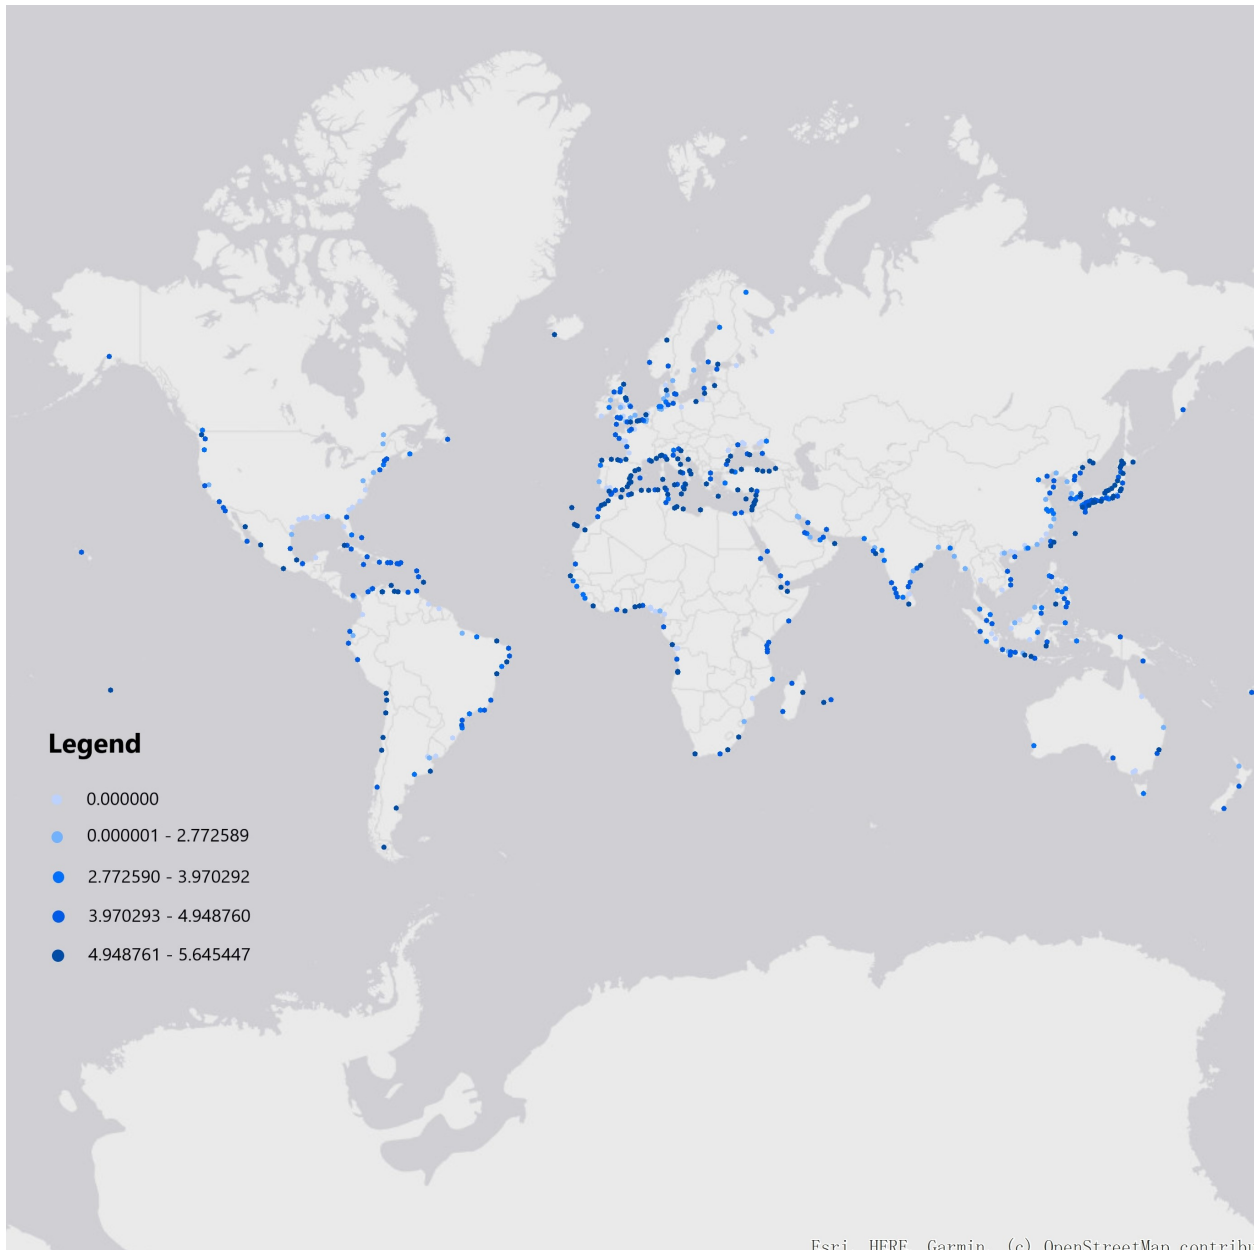


Figure 7: Variation in exogenous port suitability

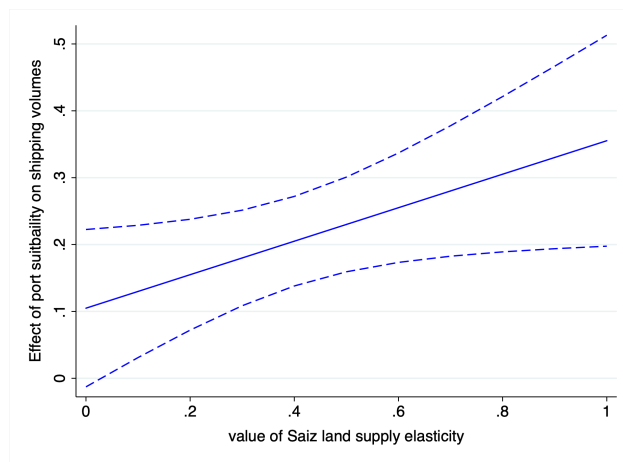


Figure 8: Shipping responds differentially to port suitability by land supply elasticity

Notes: Marginal effects from a regression of log shipping flows at the city level on port suitability interacted with a binary indicator that takes the value of 1 in decades including and after 1970 and the interaction of this measure with the Saiz land supply elasticity. Marginal effect of port suitability by land supply elasticity (low values = low land supply elasticity).

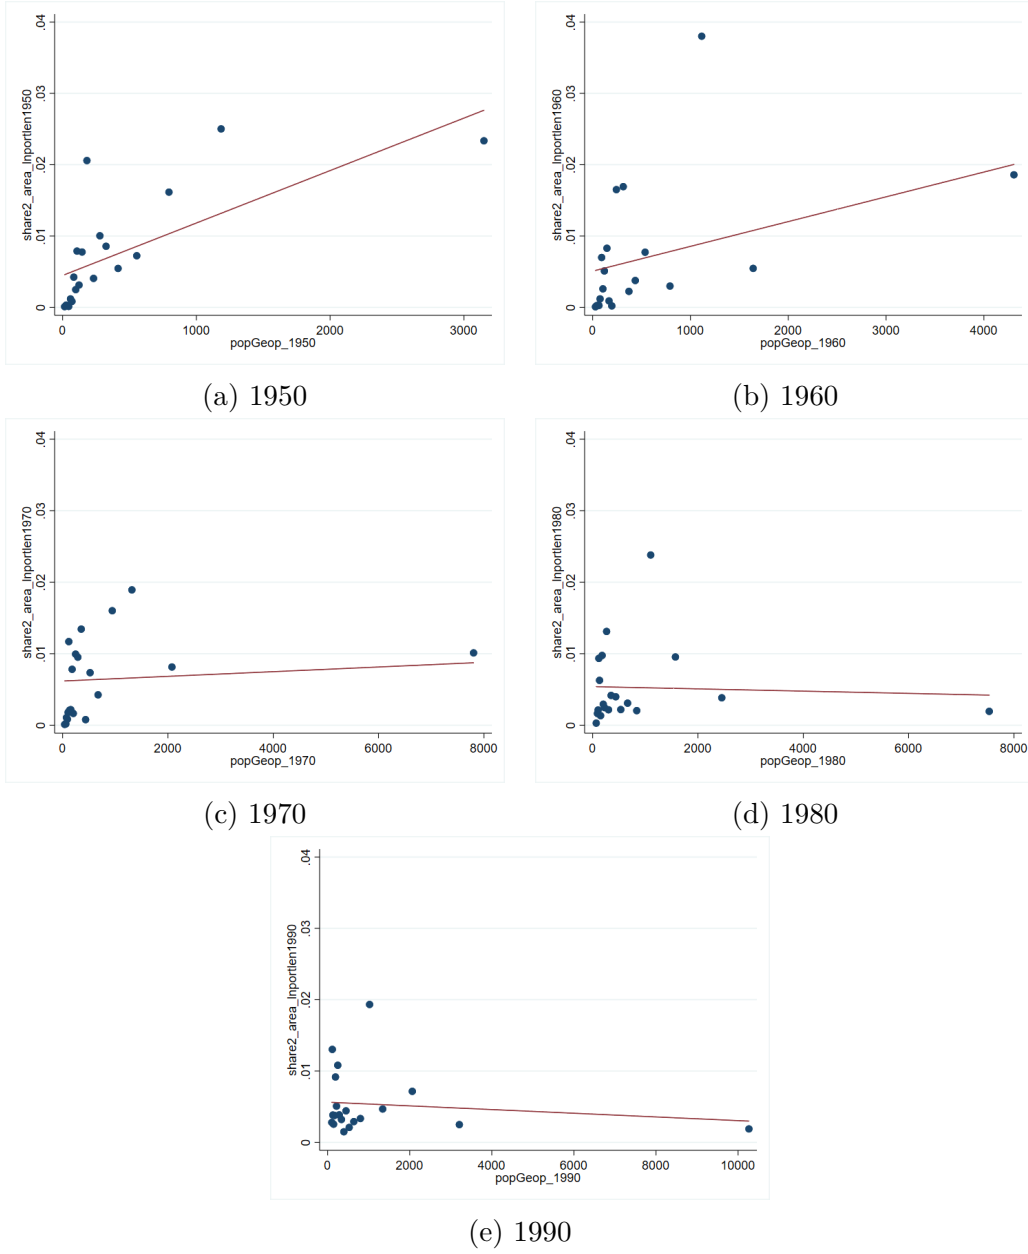


Figure 9: Correlation between port area and population, 1950-1990 (binscatter)

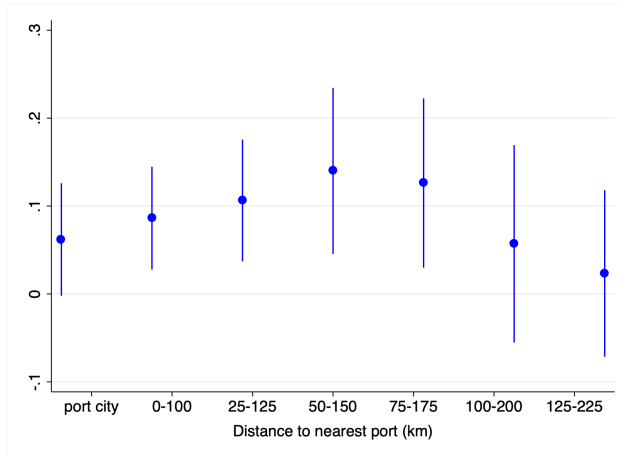


Figure 10: Estimated causal effect of shipping on population of port city and nearby cities (controlling for continent by year non-parametric time trends)
 Notes: Refer to Table 6 for details on the estimation.

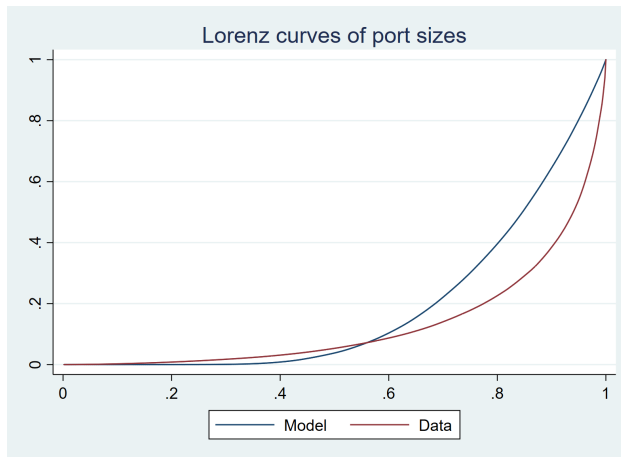


Figure 11: Lorenz curves of port shares in the model (blue) and in the data (red)




Figure 12: Change in country welfare between the counterfactual and the baseline

C Theory appendix

The primary aim of this appendix is to provide the derivation of the model's equilibrium conditions used in Section 4.2. Section C.1 derives workers' optimal location choices. Section C.2 solves the landlords' problem for the optimal allocation of land between production and transshipment. Section C.3 solves the firms' problem, while Section C.4 uses optimal prices, the price index and market clearing to obtain the equations characterizing the equilibrium distribution of wages and population. Section C.5 derives the value of shipments flowing through any port in equilibrium. Section C.6 shows how we invert the equilibrium conditions to back out amenities, productivities and exogenous port costs as a function of observed population, wages and the value of shipments. Finally, Section C.7 describes how we simulate the model in the counterfactual.

C.1 Workers' optimal location choices

In the model, a worker j who is a resident of country c and chooses to live in city r obtains utility

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

which implies that the indirect utility of a worker in city r equals

$$u_j(r) = \frac{w(r)}{P(r)} a(r) b_j(r)$$

where $w(r)$ is the nominal wage and $P(r)$ is the CES price index of consumption goods in the city.

We assume that $b_j(r)$ is distributed Fréchet with scale parameter one and shape parameter $1/\eta$:

$$Pr(b_j(r) \leq b) = e^{-b^{-1/\eta}}$$

from which we obtain that the worker's indirect utility is also distributed Fréchet with scale parameter $\frac{w(r)}{P(r)} a(r)$:

$$Pr(u_j(r) \leq u) = e^{-\left[\frac{w(r)}{P(r)} a(r)\right]^{1/\eta} u^{-1/\eta}}$$

and hence, by the properties of the Fréchet distribution, the probability with which a worker chooses to live in city r is given by

$$Pr(u_j(r) \geq u_j(s) \quad \forall s \neq r) = \frac{\left[\frac{w(r)}{P(r)} a(r)\right]^{1/\eta}}{\sum_{s \in c} \left[\frac{w(s)}{P(s)} a(s)\right]^{1/\eta}}.$$

In equilibrium, the fraction of workers choosing to live in city r coincides with this probability,

implying

$$\frac{N(r)}{\sum_{s \in c} N(s)} = \frac{\left[\frac{w(r)}{P(r)} a(r) \right]^{1/\eta}}{\sum_{s \in c} \left[\frac{w(s)}{P(s)} a(s) \right]^{1/\eta}}. \quad (6)$$

C.2 Landlords' optimal land use

Landlords earn income from providing transshipment services and from renting out land to firms that produce the city-specific good. They maximize their utility coming from their consumption of goods,

$$u_L(r) = \left[\sum_{s=1}^S q_L(s, r) \frac{\sigma-1}{\sigma} \right]^{\frac{\sigma}{\sigma-1}}$$

which implies that the indirect utility of the representative landlord in city r equals her nominal income divided by the price index,

$$u_L(r) = \frac{\left[O(r) - (\nu(r) + \psi(F(r))) Shipping(r)^\lambda \right] Shipping(r) + R(r)(1 - F(r))}{P(r)}$$

where $O(r)$ is the price of transshipment services in city r (taken as given by the landlord), $\nu(r)$ is the exogenous part of transshipment costs, $F(r)$ is the share of land allocated to the port, $Shipping(r)$ is the value of shipments flowing through the port, excluding the price of transshipment services (hence, total demand for transshipment services, again taken as given by the landlord), $R(r)$ is the land rent prevailing in the city, and $1 - F(r)$ is the share of land rented out to firms. That is, the first term in the numerator corresponds to the landlord's net nominal income from providing transshipment services, while the second term corresponds to her nominal income from renting out land to firms.

The landlord decides on the allocation of land, captured by the single variable $F(r)$, to maximize her utility. As she cannot influence the price index $P(r)$, this is equivalent to maximizing her nominal income:

$$\max_{F(r)} \left[O(r) - (\nu(r) + \psi(F(r))) Shipping(r)^\lambda \right] Shipping(r) + R(r)(1 - F(r))$$

The first-order condition to this maximization problem is

$$-\psi'(F(r)) Shipping(r)^{1+\lambda} - R(r) = 0$$

from which, by rearranging,

$$-\psi'(F(r)) = \frac{R(r)}{Shipping(r)^{1+\lambda}}. \quad (7)$$

C.3 Firms' problem

The representative firm operating in city r faces the production function

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

and maximizes its profits by choosing its labor and land use:

$$\max_{n(r), 1-F(r)} p(r, r) \tilde{A}(r) n(r)^\gamma (1 - F(r))^{1-\gamma} - w(r) n(r) - R(r) (1 - F(r))$$

where $p(r, r)$ is the factory gate price of the good produced by the firm.

The first-order conditions to this maximization problem imply

$$R(r) = \frac{1 - \gamma}{\gamma} \frac{w(r) N(r)}{1 - F(r)} \quad (8)$$

where we have used labor market clearing, which implies $n(r) = N(r)$. Plugging this back into the firm's cost function and production function, we obtain that the firm's marginal cost of production is equal to

$$\gamma^{-\gamma} (1 - \gamma)^{-(1-\gamma)} \tilde{A}(r)^{-1} w(r)^\gamma R(r)^{1-\gamma}$$

which, by perfect competition among firms, equals the factory gate price in equilibrium:

$$p(r, r) = \gamma^{-1} A(r)^{-1} (1 - F(r))^{-(1-\gamma)} N(r)^{1-\gamma-\alpha} w(r) \quad (9)$$

where we have used (8) again, together with the fact that $\tilde{A}(r) = A(r) N(r)^\alpha$.

Finally, equation (8) also implies that total factor payments in city r equal

$$Y(r) = w(r) N(r) + R(r) (1 - F(r)) = w(r) N(r) + \frac{1 - \gamma}{\gamma} w(r) N(r) = \frac{1}{\gamma} w(r) N(r). \quad (10)$$

C.4 Equilibrium conditions

From the workers' and landlords' problems, we can derive the constant-elasticity demand for the city- r good in city s as

$$q(r, s) = p(r, s)^{-\sigma} P(s)^{\sigma-1} Y(s)$$

where $p(r, s)$ is the price paid by the consumer, which includes the shipping cost between r and s . Demand in value terms is equal to

$$p(r, s) q(r, s) = p(r, r)^{1-\sigma} P(s)^{\sigma-1} Y(s) \mathbf{E}[T(r, s)]^{1-\sigma}$$

where we have used the fact that the price is the average iceberg cost over the factory gate price,

$$p(r, s) = p(r, r) \mathbf{E}[T(r, s)].$$

Market clearing for the good produced in city r implies that total factor payments in r equal demand for the good (in value terms) in the world economy:

$$\frac{1}{\gamma} w(r) N(r) = \sum_{s=1}^S p(r, r)^{1-\sigma} P(s)^{\sigma-1} \frac{1}{\gamma} w(s) N(s) \mathbf{E}[T(r, s)]$$

where we have used equation (10) to substitute for total factor payments on both sides. Plugging (9) into this equation yields

$$\begin{aligned} w(r) N(r) &= \gamma^{\sigma-1} A(r)^{\sigma-1} (1 - F(r))^{(1-\gamma)(\sigma-1)} N(r)^{-(1-\gamma-\alpha)(\sigma-1)}. \\ &w(r)^{1-\sigma} \sum_{s=1}^S P(s)^{\sigma-1} w(s) N(s) \mathbf{E}[T(r, s)]^{1-\sigma}. \end{aligned} \quad (11)$$

The CES price index in city r takes the form

$$P(r)^{1-\sigma} = \sum_{s=1}^S p(s, r)^{1-\sigma} = \sum_{s=1}^S p(s, s)^{1-\sigma} \mathbf{E}[T(s, r)]^{1-\sigma}.$$

Plugging factory gate prices (9) into this equation yields

$$P(r)^{1-\sigma} = \gamma^{\sigma-1} \sum_{s=1}^S A(s)^{\sigma-1} (1 - F(s))^{(1-\gamma)(\sigma-1)} w(s)^{1-\sigma} N(s)^{-(1-\gamma-\alpha)(\sigma-1)} \mathbf{E}[T(s, r)]^{1-\sigma}. \quad (12)$$

Rearranging equation (6) yields the following expression for the price index:

$$P(r) = \tilde{a}(r) w(r) N(r)^{-\eta} \quad (13)$$

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

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

Plugging equation (13) into (11) yields

$$\begin{aligned} &A(r)^{1-\sigma} (1 - F(r))^{-(1-\gamma)(\sigma-1)} w(r)^\sigma N(r)^{1+(1-\gamma-\alpha)(\sigma-1)} = \\ &\gamma^{\sigma-1} \sum_{s=1}^S \tilde{a}(s)^{\sigma-1} w(s)^\sigma N(s)^{1-\eta(\sigma-1)} \mathbf{E}[T(r, s)]^{1-\sigma} \end{aligned} \quad (14)$$

while plugging equation (13) into (12) yields

$$\begin{aligned} \tilde{a}(r)^{1-\sigma} w(r)^{1-\sigma} N(r)^{\eta(\sigma-1)} &= \gamma^{\sigma-1}. \\ \sum_{s=1}^S A(s)^{\sigma-1} (1-F(s))^{(1-\gamma)(\sigma-1)} w(s)^{1-\sigma} N(s)^{-(1-\gamma-\alpha)(\sigma-1)} \mathbf{E}[T(s,r)]^{1-\sigma} &. \end{aligned} \quad (15)$$

Note that our assumptions on trade costs guarantee symmetry and hence $\mathbf{E}[T(r,s)]^{1-\sigma} = \mathbf{E}[T(s,r)]^{1-\sigma}$. Given this, we can show that equations (14) and (15) can be simplified further. To see that this is the case, guess that wages take the form

$$w(r) = \tilde{a}(r)^{\iota_1} A(r)^{\iota_2} (1-F(r))^{\iota_3} N(r)^{\iota_4}.$$

That is, they only depend on local amenities, productivity, land available for production, and population. Inspecting equations (14) and (15), one can verify that this guess is indeed correct if

$$\iota_1 = -\frac{\sigma-1}{2\sigma-1} \quad \iota_2 = \iota_3 = (1-\gamma) \frac{\sigma-1}{2\sigma-1} \quad \iota_4 = [\eta - (1-\gamma)(1-\alpha)(\sigma-1) - 1] \frac{1}{2\sigma-1}$$

as (14) and (15) reduce to the same equation if the guess is correct with these values of ι_1 , ι_2 , ι_3 and ι_4 . Thus, wages in city r are given by

$$w(r) = \tilde{a}(r)^{-\frac{\sigma-1}{2\sigma-1}} A(r)^{\frac{\sigma-1}{2\sigma-1}} (1-F(r))^{(1-\gamma)\frac{\sigma-1}{2\sigma-1}} N(r)^{[\eta-(1-\gamma-\alpha)(\sigma-1)-1]\frac{1}{2\sigma-1}}. \quad (16)$$

Finally, plugging (16) back into either (14) or (15) gives us an equation that determines the distribution of population across cities:

$$N(r)^{[1+\eta\sigma+(1-\gamma-\alpha)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} = \gamma^{\sigma-1} \tilde{a}(r)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} A(r)^{\frac{(\sigma-1)^2}{2\sigma-1}} (1-F(r))^{(1-\gamma)\frac{(\sigma-1)^2}{2\sigma-1}} MA(r) \quad (17)$$

where

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

is the market access of city r .

C.5 Equilibrium shipping flows

This section derives the equilibrium value of shipping flows through any port. To obtain these, we first need to introduce further notation. Let Z be an $S+P$ by $S+P$ matrix, where P is the number of ports in the model.⁴³ Each of the first S rows and columns of Z

⁴²We can freely choose the intercept of this equation as we have not normalized any price yet. We choose it to be equal to one.

⁴³Recall that S is the total number of (port or inland) cities.

corresponds to a city, while each of the last P rows and columns of Z corresponds to a port. Let us call a city or a port a *location*; that is, each row and column in Z corresponds to one location. We assume that an entry $z(i, \ell)$ of Z is zero if locations i and ℓ are not directly connected, or if $i = \ell$. Otherwise, $z(i, \ell)$ is defined as

$$z(i, \ell) = [\bar{T}(i, \ell) [1 + O(\ell)]]^{-\theta}$$

where $\bar{T}(i, \ell)$ is the common cost of shipping from i to ℓ directly, and $O(\ell)$ is the port cost at ℓ . If ℓ is not a port but a (port or inland) city, then we define $O(\ell) = 0$.

Following Allen and Arkolakis (2019), we can show that the expected cost of shipping from city r to s can be written as

$$\mathbf{E}[T(r, s)] = \Gamma\left(\frac{\theta + 1}{\theta}\right) x(r, s)^{-1/\theta}$$

where $x(r, s)$ is the (r, s) entry of the matrix

$$X = (I - Z)^{-1}$$

and I is the $S + P$ by $S + P$ identity matrix.

Similarly, we can show that, if a good is shipped from city r to s , the probability that it is shipped through port k is given by

$$\pi(k|r, s) = \frac{x(r, k) x(k, s)}{x(r, s)} \quad (18)$$

and the probability that it is shipped through the direct link between ports k and m is

$$\pi(k, m|r, s) = \frac{x(r, k) [\bar{T}(k, m) [1 + O(m)]]^{-\theta} x(m, s)}{x(r, s)}. \quad (19)$$

The total value of goods shipped through port k from city r to city s (excluding the price paid for transshipment services at k) equals

$$\text{Shipping}(k|r, s) = [1 + O(k)]^{-1} p(r, s)^{1-\sigma} P(s)^{\sigma-1} \frac{1}{\gamma} w(s) N(s) \pi(k|r, s).$$

Combining this with $p(r, s) = p(r, r) \mathbf{E}[T(r, s)]$ as well as equations (9), (13) and (18), yields

$$\text{Shipping}(k|r, s) = \gamma^{\sigma-2} [1 + O(k)]^{-1} A(r)^{\sigma-1} (1 - F(r))^{(1-\gamma)(\sigma-1)} N(r)^{-(1-\alpha-\gamma)(\sigma-1)}.$$

$$w(r)^{1-\sigma} \tilde{a}(s)^{\sigma-1} N(s)^{1-\eta(\sigma-1)} w(s)^\sigma \mathbf{E}[T(r, s)]^{1-\sigma} \frac{x(r, k) x(k, s)}{x(r, s)}$$

and therefore the total value of shipping through port k is given by

$$Shipping(k) = \gamma^{\sigma-2} [1 + O(k)]^{-1} \sum_r D_1(r) x(r, k) \sum_s D_2(s) \frac{\mathbf{E}[T(r, s)]^{1-\sigma}}{x(r, s)} x(k, s) \quad (20)$$

where

$$D_1(r) = A(r)^{\sigma-1} (1 - F(r))^{(1-\gamma)(\sigma-1)} N(r)^{-(1-\alpha-\gamma)(\sigma-1)} w(r)^{1-\sigma}$$

and

$$D_2(s) = \tilde{a}(s)^{\sigma-1} N(s)^{1-\eta(\sigma-1)} w(s)^\sigma.$$

C.6 Inverting the model

This section describes how we invert the equilibrium conditions of the model to back out amenities, productivities and exogenous port costs as a function of observed population, wages and the value of shipments. As a first step, we use the observed data to back out port shares in the model. To this end, we combine equations (7) and (8) to obtain port shares as a function of wages $w(r)$, population $N(r)$ and the value of shipments $Shipping(r)$ in each port city r :

$$-\psi'(F(r))(1 - F(r)) = \frac{1 - \gamma}{\gamma} \frac{w(r) N(r)}{Shipping(r)^{1+\lambda}} \quad (21)$$

Given the assumptions we made on ψ' , the left-hand side of equation (21) is strictly decreasing in $F(r)$. Moreover, the left-hand side takes every real value between zero and infinity as ψ' is continuous, $\lim_{F \rightarrow 1} \psi'(F) = 0$ and $\lim_{F \rightarrow 0} \psi'(F) = -\infty$. This guarantees that solving equation (21) identifies a unique value of $F(r) \in (0, 1)$ for every port city.

The second step consists of solving for $\tilde{a}(r)$, $A(r)$ and $\nu(r)$ for the observed $N(r)$, $w(r)$ and $Shipping(r)$, as well as the $F(r)$ recovered in the previous step. This is done using an algorithm that consists of an outer loop and an inner loop. In the inner loop, we obtain the values of $\tilde{a}(r)$ that solve the system of equations

$$\tilde{a}(r)^{1-\sigma} w(r)^{1-\sigma} N(r)^{\eta(\sigma-1)} = \gamma^{\sigma-1} \sum_{s=1}^S \tilde{a}(s)^{\sigma-1} w(s)^\sigma N(s)^{1-\eta(\sigma-1)} \mathbf{E}[T(r, s)]^{1-\sigma}$$

derived from equations (14) and (15) for a *fixed* set of exogenous transshipment costs $\nu(r)$, and hence for fixed $\mathbf{E}[T(r, s)]$. For any $\mathbf{E}[T(r, s)]$, this system yields a unique solution for $\tilde{a}(r)$. Rearranging equation (16), we can then uniquely express productivity $A(r)$ as a function of the recovered $\tilde{a}(r)$:

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

In the outer loop, we search for the set of $\nu(r)$ for which the value of shipments implied by equation (20) – hence, by $N(r)$, $w(r)$, $F(r)$ and the recovered $\tilde{a}(r)$ and $A(r)$ – rationalize the shipping flows observed in the data. In practice, we start from a uniform guess of $\nu(r) = \bar{\nu}$,

then perform a large number of iterations in which we update $\nu(r)$ gradually to get closer to satisfying equation (20). We also update $\mathbf{E}[T(r, s)]$ in every iteration step. Even though we cannot prove that this procedure identifies a unique set of $\nu(r)$, the algorithm has been converging to the same fixed point for various different initial guesses on $\nu(r)$, even when guessing non-uniform distributions of $\nu(r)$ initially.

C.7 Counterfactual simulation

This section describes how we perform counterfactual simulations with the model. First, we need to choose the absolute level of amenities $a(r)$ in each city r , as the inversion only identifies amenities up to a country-level scale, $\tilde{a}(r) = \aleph_c a(r)$. Unfortunately, nothing in the data guides us with this choice. Hence, we make the simplest possible assumption by assuming that average amenities are the same across countries and are equal to one:

$$\frac{1}{C_c} \sum_{r \in c} a(r) = \frac{1}{C_c} \sum_{r \in c} \frac{\tilde{a}(r)}{\aleph_c} = 1$$

where C_c denotes the number of cities in country c . Rearranging yields

$$\aleph_c = \frac{1}{C_c} \sum_{r \in c} \tilde{a}(r)$$

and hence we can obtain the absolute level of amenities in each city r as

$$a(r) = \frac{\tilde{a}(r)}{\aleph_c} = \frac{C_c}{\sum_{s \in c} \tilde{a}(s)} \tilde{a}(r).$$

Second, we solve for the counterfactual equilibrium of the model using an algorithm that consists of three loops embedded in each other. In the innermost loop, we obtain the distribution of population $N(r)$ that solves equation (17) for a *fixed* set of \aleph_c , $F(r)$ and $Shipping(r)$ (implying that $\mathbf{E}[T(r, s)]$ are also fixed). For any \aleph_c , $F(r)$ and $Shipping(r)$, equation (17) can be shown to have a unique positive solution if

$$\alpha < 1 - \gamma + \eta$$

which holds under the values of structural parameters chosen in the calibration. Moreover, the solution can be obtained by simply iterating on equation (17), starting from any initial guess on $N(r)$. The proof of these results follows directly from the proof of equilibrium uniqueness in Allen and Arkolakis (2014).


In the middle loop, we solve for the set of country-specific \aleph_c that guarantee that the sum of city populations equals total country population in each country:

$$\sum_{r \in c} N(r) = N_c$$


where N_c denotes the exogenously given population of country c . We also solve for wages using equation (16) and for rents using equation (8).

In the outermost loop, we iterate on the distribution of port shares and shipping flows that satisfy both equations (7) and (20), also updating $\mathbf{E}[T(r, s)]$ in every step. We use the distributions of port share and shipping obtained in the inversion as our initial guesses. Even though we cannot prove that this procedure yields a unique equilibrium, we have been converging to the same distribution of endogenous variables for different initial guesses as well.

D Online Appendix: Figures



(a) Sample: ports with observed port length



(b) Sample: 14 ports with observed total port area

Figure 13: Relationship between port area and population, 1990

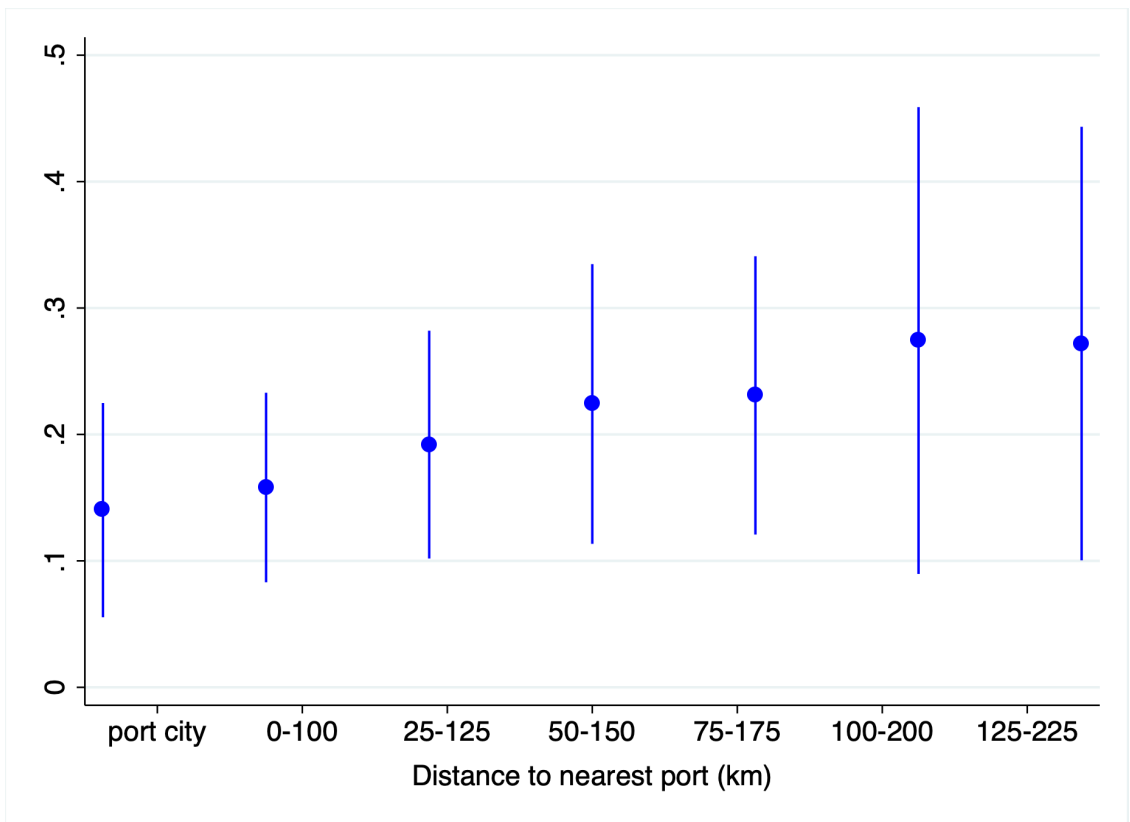


Figure 14: Estimated causal effect of shipping on population of port city and nearby cities

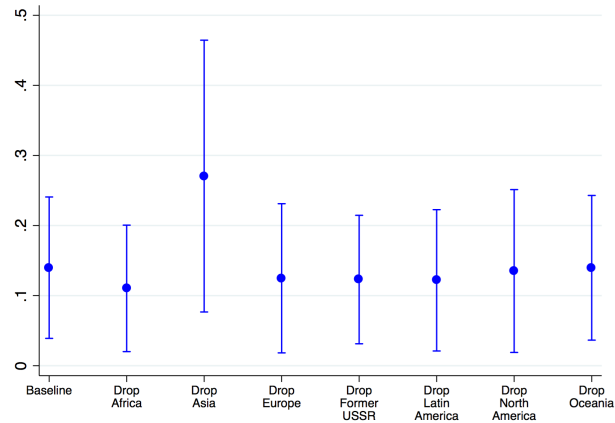

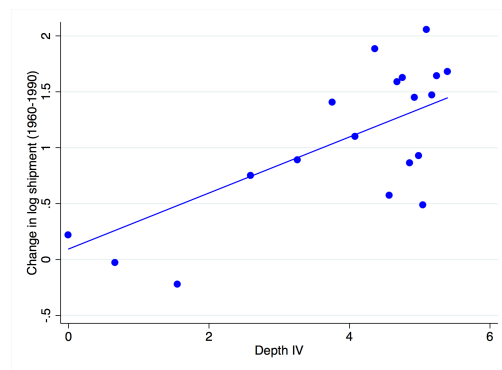


Figure 15: Reduced form effect of shipping on economic activity: Robustness to dropping continents one at a time



(a) Reduced form



(b) First stage

Figure 16: Reduced form effect of shipping on economic activity: Identifying variation in long difference (1960-1990)

E Online Appendix: Tables

Table 11: Weakening correlation between population and shipping from 1960s — robustness

Dep var: ln(shipment)	(1)	(2)	(3)	(4)
Definition	incl. non-WPI	incl. non-WPI	incl. close	incl. close
of shipment:	ports	ports	ports	ports
ln(pop)*1950	-0.075 (0.068)	-0.075 (0.076)	-0.069 (0.067)	-0.034 (0.077)
ln(pop)	1.070*** (0.079)	1.406*** (0.094)	1.105*** (0.077)	1.393*** (0.094)
ln(pop)*1970	-0.017 (0.067)	-0.154** (0.078)	-0.018 (0.065)	-0.124* (0.074)
ln(pop)*1980	-0.007 (0.080)	-0.126 (0.085)	-0.023 (0.074)	-0.099 (0.084)
ln(pop)*1990	-0.144 (0.088)	-0.308*** (0.095)	-0.152* (0.084)	-0.284*** (0.093)
Observations	2,684	2,684	2,684	2,684
R-squared	0.200	0.533	0.226	0.563
FE	year	ctryyr	year	ctryyr

Notes: Dependent variable in columns (1) and (2): log of shipping flows at the level of the city, including shipping flows from ports that do not appear in World Port Index. Dependent variable in columns (3) and (4): log of shipping flows at the level of the city, including shipping flows from ports that do not appear in World Port Index, and including shipping flows from nearby ports that are not cities themselves, and mainly serving the port city. Regressor: log of population at the level of the city interacted by a decade dummy as indicated. Statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 12: Predictive power of port suitability for shipping

depar	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)
Depth Measure X (>= 1970)	1.079*** (0.216)	0.089 (0.130)	0.199*** (0.069)	0.262*** (0.057)	0.217*** (0.041)	0.208*** (0.042)	0.172*** (0.045)	0.143*** (0.044)	0.094* (0.049)
Observations	2,609	2,609	2,609	2,609	2,609	2,609	2,609	2,609	2,609
R-squared	0.124	0.109	0.115	0.123	0.123	0.121	0.117	0.114	0.111
Number of cityid	527	527	527	527	527	527	527	527	527
year FE	yes	yes	yes	yes	yes	yes	yes	yes	yes
port FE	yes	yes	yes	yes	yes	yes	yes	yes	yes
Type	FS	FS	FS	FS	FS	FS	FS	FS	FS
Suitability measure	share	buffer 0-1km	buffer 1-3km	buffer 3-5km	buffer 5-10km	buffer 10-15km	buffer 15-20km	buffer 20-25km	buffer 25-30km

Notes: Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 13: Estimated causal effect of shipping on population of port city: robustness to the inclusion of non-parametric time trends

depvar	(1) lnpop	(2) lnpop	(3) lnpop	(4) lnpop
ln(shipment)	0.140*** (0.051)	0.062 (0.039)	0.078* (0.045)	0.189*** (0.069)
Observations	2,609	2,609	2,609	2,609
R-squared	0.505	0.732	0.662	0.428
Number of cityid	527	527	527	527
year FE	yes	yes	yes	yes
port FE	yes	yes	yes	yes
Type	2SLS	2SLS	2SLS	2SLS
FE	year	continentXyear	sizeXyear	oceanXyear
KP	27.85	24.48	25.96	19.28

Notes: Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 14: Estimated causal effect of shipping on population of port city: long difference (1960-1990)

depvar	(1) dln(population)	(2) dln(population)	(3) dln(shipment)	(4) dln(population)
dln(shipment)	0.027** (0.012)	0.141** (0.060)		
Depth Measure X (>= 1970)			0.251*** (0.059)	0.035** (0.014)
Observations	520	520	520	520
R-squared	0.012	-0.212	0.027	0.009
Type	OLS	2SLS	FS	RF
KP		18.03		

Notes: Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 15: Estimated causal effect of shipping on population of port city and nearby cities

Panel A: 2SLS estimates							
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	lngeop_pop	lngeop_pop	lngeop_pop	lngeop_pop	lngeop_pop	lngeop_pop	lngeop_pop
ln(shipment)	0.140*** (0.052)	0.158*** (0.046)	0.192*** (0.055)	0.224*** (0.067)	0.231*** (0.067)	0.274** (0.112)	0.272*** (0.104)
Observations	2,609	2,704	2,934	2,580	2,254	1,851	1,624
R-squared	0.504	0.234	0.112	-0.007	0.076	-0.056	-0.016
Number of geopolis_id	527	547	596	524	460	381	336
sample	portcity	0 - 100	25 - 125	50 - 150	75 - 175	100 - 200	125 - 225
KP	27.82	32.96	27.08	21.66	21.76	9.774	11.86
Panel B: Dynamic first stage							
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)	ln(shipment)
ivX1960	-0.007 (0.046)	-0.038 (0.055)	-0.072 (0.052)	-0.063 (0.055)	-0.017 (0.057)	-0.095 (0.059)	-0.150** (0.063)
ivX1970	0.232*** (0.049)	0.273*** (0.063)	0.251*** (0.060)	0.250*** (0.064)	0.250*** (0.062)	0.105 (0.068)	0.063 (0.068)
ivX1980	0.163*** (0.059)	0.192** (0.075)	0.118* (0.071)	0.125* (0.075)	0.140* (0.073)	0.067 (0.073)	0.086 (0.076)
ivX1990	0.246*** (0.065)	0.451*** (0.082)	0.356*** (0.077)	0.314*** (0.079)	0.315*** (0.076)	0.196** (0.077)	0.186** (0.081)
Observations	2,609	2,704	2,934	2,580	2,254	1,851	1,624
R-squared	0.125	0.123	0.130	0.133	0.138	0.159	0.158
Number of geopolis_id	527	547	596	524	460	381	336
sample	portcity	0 - 100	25 - 125	50 - 150	75 - 175	100 - 200	125 - 225
Panel C: Dynamic reduced form							
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	lngeop_pop	lngeop_pop	lngeop_pop	lngeop_pop	lngeop_pop	lngeop_pop	lngeop_pop
ivX1960	0.007 (0.006)	0.018*** (0.006)	0.017*** (0.006)	0.019*** (0.006)	0.016** (0.006)	0.015** (0.007)	0.020** (0.008)
ivX1970	0.027*** (0.009)	0.046*** (0.010)	0.047*** (0.009)	0.051*** (0.011)	0.051*** (0.011)	0.044*** (0.013)	0.050*** (0.014)
ivX1980	0.032** (0.013)	0.058*** (0.014)	0.061*** (0.013)	0.065*** (0.014)	0.058*** (0.014)	0.046*** (0.016)	0.054*** (0.017)
ivX1990	0.042** (0.016)	0.077*** (0.017)	0.079*** (0.016)	0.088*** (0.017)	0.083*** (0.017)	0.074*** (0.019)	0.080*** (0.021)
Observations	2,609	2,704	2,934	2,580	2,254	1,851	1,624
R-squared	0.627	0.648	0.658	0.678	0.690	0.693	0.707
Number of geopolis_id	527	547	596	524	460	381	336
sample	portcity	0 - 100	25 - 125	50 - 150	75 - 175	100 - 200	125 - 225

Notes: Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 16: Model-based specification: robustness to the inclusion of non-parametric time trends

depvar	(1) ln(pop)	(2) ln(pop)	(3) ln(pop)	(4) ln(pop)
ln(shipment)	-0.251*** (0.070)	-0.223*** (0.056)	-0.261*** (0.091)	-0.314*** (0.084)
ln(market access)	17.797*** (2.640)	14.291*** (2.195)	17.204*** (6.076)	24.954*** (6.294)
Observations	2,571	2,571	2,571	2,571
Number of cityid	519	519	519	519
year FE	yes	yes	yes	yes
port FE	yes	yes	yes	yes
Type	2SLS	2SLS	2SLS	2SLS
FE	year	sizemyr	continentyr	oceanyr
KP	14.27	13.19	4.035	5.943

Notes: Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 17: Model-based specification: robustness to dropping cities close to the port city in the market access IV

depvar	(1) ln(pop)	(2) ln(pop)	(3) ln(pop)	(4) ln(pop)	(5) ln(pop)	(6) ln(pop)
ln(shipment)	-0.251*** (0.070)	-0.232*** (0.068)	-0.215*** (0.067)	-0.198*** (0.068)	-0.177** (0.070)	-0.096 (0.069)
ln(market access)	17.797*** (2.640)	16.906*** (2.416)	16.154*** (2.199)	15.391*** (1.989)	14.434*** (1.763)	10.758*** (1.381)
Observations	2,571	2,571	2,571	2,571	2,571	2,571
R-squared	-1.346	-1.124	-0.947	-0.778	-0.579	0.038
Number of cityid	519	519	519	519	519	519
year FE	yes	yes	yes	yes	yes	yes
port FE	yes	yes	yes	yes	yes	yes
Type	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
FE	year	year	year	year	year	year
Drop cities in market access IV	none	≤ 100	≤ 200	≤ 300	≤ 500	≤ 1000
KP	14.27	14.79	15.52	15.99	16.09	12.26

Notes: Standard errors clustered at the city level. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.