Online Appendix

All aboard: The effects of port development

César Ducruet   Réka Juhász   Dávid Krisztián Nagy   Claudia Steinwender

A Historical evidence: The cost-space trade-off in port development

Adapting ports to containerized technology required a substantial increase in the space devoted to the port. The faster turnaround times which reduced transshipment costs could only be achieved by building much larger terminals. This space-cost trade-off was not a novel aspect of containerization. Prior developments in port technology followed a similar pattern. To understand the land-intensity of containerization, it is helpful to review the factors that drive this trade-off in more detail.

Historically, breakbulk ports had used narrow "finger-piers" which made it possible for many ships to dock at the port simultaneously. With slow loading and unloading times, ports needed to be able to accommodate many ships at the same time.

Even before containerization, small continuous technological improvements such as the use of cranes, lift machines, tractors, trailers and belts led to a decrease in loading and unloading times (Port of New Orleans, 1951, p. 22). These improvements highlighted an important trade-off between space and speed. If cargo was unloaded faster, it piled up at the wharf. Increases in turnaround times could only be realized by allocating more space at the pier. As a result, modernizing ports were developing new wharf designs to realize the gains from increased turnaround times. For example, the New York Port Authority invested in reconstructing piers at its Hoboken, Brooklyn and Newark terminals with much wider piers allowing for more efficient turnaround times. In 1955, the Port Authority describes the most modern wharf designs as follows; "large cargo terminals on one level, adjacent to shipping berths, [that] increase efficiency of freight handling, speed up ship turnaround, and permit prompt loading and discharge of trucks. *The resulting savings more than compensate for the added costs of space*" (1955, own emphasis, p.5). At the Hoboken piers, operations became 25% more efficient following the reconstruction of piers into the wider format (New York Port Authority, 1955). Similarly, contemporaries realized the advantages of single-storey warehouses over multi-storey warehouses for faster access to stored cargo and started developing large "upland areas" – that is, areas away from the terminal to make transshipment more efficient (New York Port Authority, 1955, p.7).

Following the logic above, containerized terminals need more space as it is the easy accessibility of the containers that allows for efficient on- and off-loading. The containers are lined up next to where the ships dock, and space is also needed to rapidly off-load cargo. There are additional dedicated 'upland areas' near the facility that allow for the containers to be temporarily stored

(New York Port Authority, 1958, p.5) and new space needed to be made for large 'railyards' where containers could await transshipment onto rail carriages (Riffenburgh, 2012, pp. xi-xii).[1]

The increased space requirements of containerized facilities were evident from the earliest days of the new technology. Before containerization, there were some smaller innovations in handling breakbulk cargo such as palletization (whereby goods are placed on a pallet and handled as a unit) and pre-slinging (whereby goods are grouped together using slings and the unit is handled together).[2] These new processes also allowed cargo to be handled as a unit, saving some transshipment time, though importantly, *not* in a standardized way (given the different-sized cargo involved).

However, these pre-containerization changes were minor compared to the effect of containerization. As early as 1958 (two years after the first containerized shipments had sailed from New York), the New York Port Authority put in place plans to develop the Elizabeth facility for containerized cargo handling; "Extensive supporting upland area is one of the most important features of the development, since these large open spaces are indispensable in the handling of general cargo in the age of container ships" (1958, p. 5). The Port of San Francisco (the fifth largest port in the U.S. in 1950 according to our data) was lamenting the inadequacy of the city's finger piers to accommodate new types of cargo handling; "The Port [should] commence the phasing out of finger piers. [The piers are] commercially obsolete for the new generation of ships and the new types of cargo handling technology" (Port of San Francisco, 1971, p. 27).

B Theory

B.1 Equilibrium of the model

We define the equilibrium of the model as follows.

Definition. *Given structural parameters α, γ, η, σ, ρ, θ, λ, the number of cities S and the subset of port cities $P \subseteq \{1, ..., S\}$, country populations N_c, city amenities $a : \{1, ..., S\} \to \mathbb{R}$, productivities $A : \{1, ..., S\} \to \mathbb{R}$, exogenous transshipment costs $\nu : P \to \mathbb{R}$, inland and sea shipping costs as a function of distance $\phi_\varsigma, \phi_\tau : \mathbb{R} \to \mathbb{R}$ and endogenous transshipment costs as a function of port share $\psi : (0, 1) \to \mathbb{R}$, an **equilibrium** of the model is a set of city populations $N : S \to \mathbb{R}$, nominal wages $w : S \to \mathbb{R}$, land rents $R : S \to \mathbb{R}$, employment levels $n : S \to \mathbb{R}$, port shares $F : S \to [0, 1)$, port-level shipping flows $Shipping : P \to \mathbb{R}$, the prices of transshipment services $O : P \to \mathbb{R}$, the prices of goods $p : S^2 \to \mathbb{R}$ and the quantities of goods $q : S^2 \to \mathbb{R}$ such that*

1. *workers choose their consumption of goods and city of residence within their country to*

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

[2] UNCTAD (1971) gives a more detailed overview of these trends.

maximize their utility (4), taking prices and wages as given;

2. *landlords in each city r choose their consumption of goods and land use to maximize their utility*

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

taking prices, land rents and shipping flows as given;[3]

3. *competition among landlords drives the price of transshipment services down to marginal cost, (6), and landlords' profits from transshipment down to zero;*[4]

4. *firms in each city r choose their production, employment and land use to maximize their profits,*

$$\max_{n(r),1-F(r)} p(r,r)\,\tilde{A}(r)\,n(r)^{\gamma}(1-F(r))^{1-\gamma} - w(r)\,n(r) - R(r)(1-F(r)) \tag{B.2}$$

taking prices, land rents and wages as given, where $p(r,r)$ is the factory gate price of the good produced by the firm, and choose the shipping route to each destination to maximize their profits;

5. *competition among firms drives their profits down to zero;*

6. *there is no possibility of arbitrage, implying that the price of good r at s equals the expected iceberg cost over the factory gate price,*

$$p(r,s) = p(r,r)\,\mathbf{E}\left[T(r,s)\right]; \tag{B.3}$$

7. *the market for labor clears in each city r, implying $n(r) = N(r)$;*

8. *national labor markets clear, implying $\sum_{r\in c} N(r) = N_c$ in each country c;*

9. *the market for land clears in each city;*

10. *the market for transshipment services clears in each port city;*

11. *the market for each good clears worldwide.*

[3] We assume that landlords do not enjoy city amenities and do not have idiosyncratic tastes for cities. As landlords are immobile, this assumption does not have any consequence on their optimal choices and is therefore without loss of generality.

[4] We relax this assumption in the monopolistic competition version of the model, presented in Appendix B.7.

Note that this 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.[5]

B.2 Equilibrium land use, wages, city populations and shipping flows

This section uses the equilibrium conditions of Appendix B.1 to characterize cities' equilibrium land use, wages, populations and shipping flows. To obtain these, we proceed as follows. Appendix B.2.1 solves for workers' optimal location choices. Appendix B.2.2 solves the landlords' problem for the optimal allocation of land between production and transshipment. Appendix B.2.3 solves the firms' problem, while Appendix B.2.4 uses equilibrium prices, the price index and market clearing to obtain the equations characterizing cities' equilibrium wages and population. Finally, Appendix B.2.5 derives the value of shipments flowing through any port in equilibrium.

B.2.1 Workers' optimal location choices

The utility function of workers, (4), implies that the indirect utility of a worker living 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 $\left[\frac{w(r)}{P(r)} a(r)\right]^{1/\eta}$:

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

and hence, by the properties of the Fréchet distribution, the probability with which a worker

[5]In the monopolistic competition version of the model (Appendix B.7), we do not need to make this assumption. In that model, landlords have market power and therefore choose both the price and the quantity of transshipment in a way that maximizes their profits.

chooses to live in city r is given by

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

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

$$\frac{N\left(r\right)}{\sum_{s\in c} N\left(s\right)} = \frac{\left[\frac{w(r)}{P(r)}a\left(r\right)\right]^{1/\eta}}{\sum_{s\in c}\left[\frac{w(s)}{P(s)}a\left(s\right)\right]^{1/\eta}}. \tag{B.4}$$

B.2.2 Landlords' optimal land use

Landlords earn income from providing transshipment services and from renting out land to firms that produce the city-specific good. Their utility function, (B.1), implies that the indirect utility of a landlord in city r equals her nominal income divided by the price index,

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

where $O\left(r\right)$ is the price of transshipment services in city r (taken as given by the landlord), $\nu\left(r\right)$ is the exogenous part of transshipment costs, $F\left(r\right)$ is the share of land allocated to the port, $Shipping\left(r\right)$ 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), $R\left(r\right)$ is the land rent prevailing in the city, and $1 - F\left(r\right)$ 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 $F\left(r\right)$, to maximize her utility. As she cannot influence the price index $P\left(r\right)$, this is equivalent to maximizing her nominal income:

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

The first-order condition to this maximization problem is

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

from which, by rearranging,

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

B.2.3 Firms' problem

Recall that the representative firm operating in city r faces the production function

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

and maximizes its profits, (B.2), by choosing its employment and land use. The first-order conditions to the firm's profit-maximization problem imply

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

where we have used labor market clearing, which implies $n\left(r\right)=N\left(r\right)$. 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

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

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

$$p\left(r,r\right)=\gamma^{-1}A\left(r\right)^{-1}\left(1-F\left(r\right)\right)^{-(1-\gamma)}N\left(r\right)^{1-\gamma-\alpha}w\left(r\right) \tag{B.7}$$

where we have used (B.6) again, together with the fact that $\tilde{A}\left(r\right)=A\left(r\right)N\left(r\right)^{\alpha}$.

Finally, equation (B.6) also implies that total factor payments in city r equal

$$Y\left(r\right)=w\left(r\right)N\left(r\right)+R\left(r\right)\left(1-F\left(r\right)\right)=w\left(r\right)N\left(r\right)+\frac{1-\gamma}{\gamma}w\left(r\right)N\left(r\right)=\frac{1}{\gamma}w\left(r\right)N\left(r\right). \tag{B.8}$$

B.2.4 Equilibrium wages and populations

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

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

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

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

where we have used equation (B.3).

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

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

where we have used equation (B.8) to substitute for total factor payments on both sides. Plugging (B.7) into this equation yields

$$\begin{aligned} w(r)\,N(r)=&\gamma^{\sigma-1}A(r)^{\sigma-1}\left(1-F(r)\right)^{(1-\gamma)(\sigma-1)}N(r)^{-(1-\gamma-\alpha)(\sigma-1)}\cdot\\ &w(r)^{1-\sigma}\sum_{s=1}^{S}P(s)^{\sigma-1}\,w(s)\,N(s)\,\mathbf{E}\left[T(r,s)\right]^{1-\sigma}. \end{aligned} \tag{B.9}$$

The CES price index in city r takes the form

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

Plugging factory gate prices (B.7) into this equation yields

$$P(r)^{1-\sigma}=\gamma^{\sigma-1}\sum_{s=1}^{S}A(s)^{\sigma-1}\left(1-F(s)\right)^{(1-\gamma)(\sigma-1)}w(s)^{1-\sigma}\,N(s)^{-(1-\gamma-\alpha)(\sigma-1)}\,\mathbf{E}\left[T(s,r)\right]^{1-\sigma}. \tag{B.10}$$

Rearranging equation (B.4) yields the following expression for the price index:

$$P(r)=\tilde{a}(r)\,w(r)\,N(r)^{-\eta} \tag{B.11}$$

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 (B.11) into (B.9) yields

$$\begin{aligned} &A(r)^{1-\sigma}\left(1-F(r)\right)^{-(1-\gamma)(\sigma-1)}w(r)^{\sigma}\,N(r)^{1+(1-\gamma-\alpha)(\sigma-1)}=\\ &\gamma^{\sigma-1}\sum_{s=1}^{S}\tilde{a}(s)^{\sigma-1}\,w(s)^{\sigma}\,N(s)^{1-\eta(\sigma-1)}\,\mathbf{E}\left[T(r,s)\right]^{1-\sigma} \end{aligned} \tag{B.12}$$

while plugging equation (B.11) into (B.10) yields

$$\tilde{a}(r)^{1-\sigma} w(r)^{1-\sigma} N(r)^{\eta(\sigma-1)} = \gamma^{\sigma-1} \cdot \sum_{s=1}^{S} A(s)^{\sigma-1} (1-F(s))^{(1-\gamma)(\sigma-1)} w(s)^{1-\sigma} N(s)^{-(1-\gamma-\alpha)(\sigma-1)} \mathbf{E}[T(s,r)]^{1-\sigma}. \tag{B.13}$$

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 (B.12) and (B.13) 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 (B.12) and (B.13), one can verify that this guess is indeed correct if

$$\iota_1 = -\frac{\sigma-1}{2\sigma-1},$$

$$\iota_2 = \iota_3 = (1-\gamma)\frac{\sigma-1}{2\sigma-1}$$

and

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

as (B.12) and (B.13) 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

$$w(r) = \tilde{a}(r)^{-\frac{\sigma-1}{2\sigma-1}} A(r)^{\frac{\sigma-1}{2\sigma-1}} (1-F(r))^{(1-\gamma)\frac{\sigma-1}{2\sigma-1}} N(r)^{[\eta-(1-\gamma-\alpha)(\sigma-1)-1]\frac{1}{2\sigma-1}}. \tag{B.14}$$

[6]

Finally, plugging (B.14) back into either (B.12) or (B.13) gives us an equation that determines the distribution of population across cities:

$$N(r)^{[1+\eta\sigma+(1-\gamma-\alpha)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} = \gamma^{\sigma-1}\tilde{a}(r)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} A(r)^{\frac{(\sigma-1)^2}{2\sigma-1}} (1-F(r))^{(1-\gamma)\frac{(\sigma-1)^2}{2\sigma-1}} MA(r) \tag{B.15}$$

where

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

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

is the market access of city r.

B.2.5 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 denotes both the set and the number of ports in the model.[7] Each of the first S rows and columns of Z 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

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

where $\bar{T}(i,\ell)$ is the common cost of shipping from i to ℓ directly, and $O(\ell)$ is the price of transshipment services at ℓ. If ℓ is a port belonging to port city r, then this price is given by equation (6). If ℓ is not a port but a (port or non-port) city, then we define $O(\ell)=0$.[8]

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

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

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

$$\pi(k|r,s)=\frac{x(r,k)\,x(k,s)}{x(r,s)} \tag{B.16}$$

and therefore 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

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

[7] Recall that S is the total number of (port or non-port) cities.

[8] For computational reasons, we need to add a small iceberg cost of shipping between each port and its own city. This cost equals 1.03 in both the inversion and the model simulations.

Combining this with equations (B.3), (B.7), (B.11) and (B.16) yields

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

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

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

$$Shipping\left(k\right)=\gamma^{\sigma-2}\left[1+O\left(k\right)\right]^{-1}\sum_{r}D_{1}\left(r\right)x\left(r,k\right)\sum_{s}D_{2}\left(s\right)\frac{\mathbf{E}\left[T\left(r,s\right)\right]^{1-\sigma}}{x\left(r,s\right)}x\left(k,s\right) \tag{B.17}$$

where

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

and

$$D_{2}\left(s\right)=\tilde{a}\left(s\right)^{\sigma-1}N\left(s\right)^{1-\eta\left(\sigma-1\right)}w\left(s\right)^{\sigma}.$$

B.3 Inverting the model

This section describes how we invert the equilibrium conditions of the model to back out amenities up to a country-level scale, productivities and exogenous transshipment 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 (B.5) and (B.6) to obtain port shares as a function of wages $w\left(r\right)$, population $N\left(r\right)$ and the value of shipments $Shipping\left(r\right)$ in each port city r:

$$-\psi'\left(F\left(r\right)\right)\left(1-F\left(r\right)\right)=\frac{1-\gamma}{\gamma}\frac{w\left(r\right)N\left(r\right)}{Shipping\left(r\right)^{1+\lambda}} \tag{B.18}$$

Given the assumptions we made on ψ', the left-hand side of equation (B.18) is strictly decreasing in $F\left(r\right)$. Moreover, the left-hand side takes every real value between zero and infinity as ψ' is continuous, $\lim_{F\to1}\psi'\left(F\right)=0$ and $\lim_{F\to0}\psi'\left(F\right)=-\infty$. This guarantees that solving equation (B.18) identifies a unique value of $F\left(r\right)\in\left(0,1\right)$ for every port city.

The second step consists of solving for $\tilde{a}\left(r\right)$, $A\left(r\right)$ and $\nu\left(r\right)$ for the observed $N\left(r\right)$, $w\left(r\right)$ and $Shipping\left(r\right)$, as well as the $F\left(r\right)$ 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 $\tilde{a}\left(r\right)$ that solve the system of equations

$$\tilde{a}\left(r\right)^{1-\sigma}w\left(r\right)^{1-\sigma}N\left(r\right)^{\eta\left(\sigma-1\right)}=\gamma^{\sigma-1}\sum_{s=1}^{S}\tilde{a}\left(s\right)^{\sigma-1}w\left(s\right)^{\sigma}N\left(s\right)^{1-\eta\left(\sigma-1\right)}\mathbf{E}\left[T\left(r,s\right)\right]^{1-\sigma} \tag{B.19}$$

derived from equations (B.12) and (B.13) for a *fixed* set of exogenous transshipment costs $\nu\left(r\right)$, 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 (B.14), we can then uniquely express productivity $A(r)$ as a function of the recovered $\tilde{a}(r)$:

$$A(r) = \tilde{a}(r)(1 - F(r))^{\gamma-1} w(r)^{\frac{2\sigma-1}{\sigma-1}} N(r)^{-[\eta-(1-\gamma-\alpha)(\sigma-1)-1]\frac{1}{\sigma-1}} \tag{B.20}$$

In the outer loop, we search for the set of $\nu(r)$ for which the value of shipments implied by equation (B.17) – 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 (B.17). 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.

B.4 Counterfactual simulation

This section describes how we perform counterfactual simulations in 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

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

We also need to remove endogenous, shipping-related disamenities from $a(r)$ and recover fundamental amenities $\bar{a}(r)$. This can be done by rearranging equation (5):

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

where $Shipping\left(r\right)$ is the observed value of shipments at 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 $N\left(r\right)$ that solves equation (B.15) for a *fixed* set of $\aleph_c$, $F\left(r\right)$ and $Shipping\left(r\right)$ (implying that $\mathbf{E}\left[T\left(r,s\right)\right]$ are also fixed). For any $\tilde{a}\left(r\right)$, $F\left(r\right)$ and $Shipping\left(r\right)$, equation (B.15) can be shown to have a unique positive solution if

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

which holds under the assumptions made in Section 4.1. Moreover, the solution can be obtained by simply iterating on equation (B.15), starting from any initial guess on $N\left(r\right)$. 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:

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

where N_c denotes the exogenously given population of country c. We also solve for wages using equation (B.14) and for rents using equation (B.6).

In the outermost loop, we iterate on the distribution of port shares and shipping flows that satisfy both equations (B.5) and (B.17), also updating $\mathbf{E}\left[T\left(r,s\right)\right]$ 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.

B.5 Benchmark models

This section provides a description of the two benchmark models (Benchmark 1 and Benchmark 2) used to decompose the local and aggregate effects of containerization. To implement these decompositions, we first take Benchmark 1 and Benchmark 2 to our 1990 data. Next, we conduct the no-containerization counterfactual in each benchmark model. In particular, we conduct the counterfactual such that the world trade to GDP ratio changes to the same extent ($+4.1$ pp) in each benchmark as in our baseline model. Hence, differences in the welfare effects across the models do not stem from trade changing to a different extent.

Here, we introduce each benchmark model in detail and show how they are taken to the data and how they are solved for the counterfactual (no-containerization) equilibrium.

B.5.1 Benchmark 1: No land use in transshipment

In Benchmark 1, we abstract from endogenous (land-dependent) transshipment costs. Thus, the cost of handling one unit of a good at port p_m is given by

$$\nu(p_m)\,Shipping(p_m)^{\lambda}$$

and, by perfect competition, the price of transshipment services equals this cost:

$$O(p_m) = \nu(p_m)\,Shipping(p_m)^{\lambda} \tag{B.21}$$

As production is the only sector in which land can be productively used in this model, landlords optimally set the fraction of production land to one: $1 - F(r) = 1$. The remaining assumptions are the same as in the baseline model.[9] Naturally, equation (B.5) does not hold in Benchmark 1, since all port shares are equal to zero.

Taking Benchmark 1 to the data. Taking Benchmark 1 to 1990 data follows similar steps as taking our baseline model to the data. We keep the structural parameters and the inland and sea shipping costs unchanged relative to the baseline model. To back out amenities, productivities and exogenous transshipment costs after containerization, we invert Benchmark 1 using 1990 data on population, wages and the value of shipments. This inversion procedure differs from the inversion of the baseline model in that we do not need to solve equation (B.5) for equilibrium port shares. As a result, we can skip the first step of the inversion procedure and immediately start with what we labeled as the second step in Appendix B.3.

In particular, we solve an algorithm that consists of an outer loop and an inner loop. In the inner loop, we obtain the values of city amenities $\tilde{a}(r)$ that solve equation (B.19), which holds in Benchmark 1 as well, for a *fixed* set of $\nu(r)$, hence for fixed $\mathbf{E}[T(r,s)]$. Once we have $\tilde{a}(r)$, we can obtain city productivities $A(r)$ from equation (B.20), which also holds in Benchmark 1, such that we set $1 - F(r) = 1$.

In the outer loop, we search for the set of $\nu(r)$ such that shipments implied by equation (B.17) equal the shipping flows observed in the data. Equation (B.17) also holds in Benchmark 1, except that we need to use $1 - F(r) = 1$ and equation (B.21) instead of equation (6) to calculate transshipment prices. 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 (B.17). We also update $\mathbf{E}[T(r,s)]$ in every iteration step.

[9] We keep port city disamenities in the version of Benchmark 1 that we label as 'Benchmark 1 with disamenities.' In the version without disamenities, which we use in column (3) of Table 4, we assume that the amenity level of any city r, $a(r)$, is exogenously given.

Counterfactual simulation of Benchmark 1. When conducting the no-containerization counterfactual in Benchmark 1, we again stay as close as possible to our baseline model. We offset the relationship between $\log \nu(r)$ and port depth, and increase all $\log \nu(r)$ by a constant ν_{CF} such that we have the same increase in international trade to world GDP as in the baseline model (Section 6.2). We also use the same procedure to obtain $\bar{a}(r)$ from $\tilde{a}(r)$ (Appendix B.4). When conducting the targeted port development counterfactual, we decrease exogenous transshipment costs in the 24 targeted ports by 10%, as in the baseline model.

We solve for counterfactual equilibria 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 (B.15) for a *fixed* set of $\aleph_c$ and $Shipping(r)$ (implying that $\mathbf{E}[T(r,s)]$ are also fixed). Equation (B.15) is unchanged relative to the baseline model, except that we need to use $1-F(r) = 1$. We follow the same iterative procedure as in Appendix B.4 to solve equation (B.15).

In the middle loop, we solve for the set of country-specific $\aleph_c$ such that the sum of city populations equals total country population in each country. We also solve for wages using equation (B.14), which is the same as in the baseline model, except that $1 - F(r) = 1$.

In the outermost loop, we iterate on equation (B.17) to obtain equilibrium shipping flows, also updating $\mathbf{E}[T(r,s)]$ in every step. In contrast to the baseline model, we use $1 - F(r) = 1$ and equation (B.21) instead of equation (6) in this process. We use the 1990 shipping flows as our initial guess.

B.5.2 Benchmark 2: Land use in transshipment identical across port cities

In Benchmark 2, we allow for endogenous (land-dependent) transshipment costs. This implies that transshipment prices are given by equation (6), just like in our baseline model. However, we restrict transshipment land use to be identical across port cities. More precisely, we set the 1990 port share of each port city equal to the average 1990 port share in the baseline model. Similarly, we set the no-containerization counterfactual port share equal to the average port share in the no-containerization counterfactual of our baseline model. The remaining assumptions are the same as in the baseline model. Similar to Benchmark 1, equation (B.5) does not hold in this model since port shares are set exogenously through the above procedure, rather than optimally by port city landlords.

Taking Benchmark 2 to the data. We keep the structural parameters and the inland and sea shipping costs unchanged relative to the baseline model. To back out amenities, productivities and exogenous transshipment costs after containerization, we invert Benchmark 2 using 1990 data on population, wages and the value of shipments. Just like in Benchmark 1, we do not need to solve equation (B.5) for equilibrium port shares. As a result, we can skip the first step of the inversion procedure and immediately start from the second step. This second step, in turn, is conducted exactly as in the baseline model (see Appendix B.3 for details), except that we use the average 1990 port share in the baseline model as $F(r)$ in each port city.

Counterfactual simulation of Benchmark 2. In the no-containerization counterfactual simulation of Benchmark 2, we change transshipment cost parameter β in the same way as in the counterfactual of the baseline model; offset the relationship between $\log \nu(r)$ and port depth; and increase all $\log \nu(r)$ by a constant ν_{CF} such that we have the same increase in international trade to world GDP as in the baseline model (Section 6.2). We also use the same procedure to obtain $\bar{a}(r)$ from $\tilde{a}(r)$ (Appendix B.4).

Finally, we solve for the counterfactual equilibrium 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 (B.15) for a *fixed* set of $\aleph_c$, $F(r)$ and $Shipping(r)$ (implying that $\mathbf{E}[T(r,s)]$ are also fixed). We use the average port share in the counterfactual of the baseline model as $F(r)$ in each port city. We follow the same iterative procedure as in Appendix B.4 to solve equation (B.15).

In the middle loop, we solve for the set of country-specific $\aleph_c$ such that the sum of city populations equals total country population in each country. We also solve for wages using equation (B.14), which is the same as in the baseline model. We again use the same $F(r)$ in each port city.

In the outermost loop, we iterate on equation (B.17) to obtain equilibrium shipping flows, also updating $\mathbf{E}[T(r,s)]$ in every step. We again use the same $F(r)$ in each port city. We use the 1990 shipping flows as our initial guess.

B.6 A model with labor used in transshipment

This section presents a generalization of our baseline model in which the provision of transshipment services may require not only land, but also potentially labor. We show that, as long as the share of labor relative to land in transshipment is sufficiently low, this more general framework delivers predictions on port development and city populations that are similar to the predictions of our baseline model. On the other hand, if the share of labor in transshipment is high, the model's predictions are in contrast with the empirical findings of Section 5, as we describe below.

We now present the setup of the model with transshipment labor. Assume that the cost of transshipping one unit of a good in port city r equals

$$\left(\nu(r) + \psi\left(n^P(r)^{\gamma_P} F(r)^{1-\gamma_P}\right)\right) Shipping(r)^{\lambda}$$

where $0 \leq \gamma_P \leq 1$. That is, γ_P is labor's share and $1-\gamma_P$ is land's share in transshipment services. Our baseline model is a special case in which $\gamma_P = 0$. The remaining model assumptions are the same as in the baseline model.

We now show how our model predictions – more precisely, Proposition 1 – change in this more general framework. First, note that the first-order conditions to the landlord's problem with respect to $n^P(r)$ and $F(r)$ together imply

$$n^P(r) = \frac{\gamma_P}{1-\gamma_P}\frac{R(r)}{w(r)}F(r). \tag{B.22}$$

On the production side, the first-order conditions to the firm's problem imply

$$n(r) = \frac{\gamma}{1-\gamma}\frac{R(r)}{w(r)}(1-F(r)). \tag{B.23}$$

Adding equations (B.22) and (B.23) yields total demand for labor in the city,

$$N(r) = \frac{\gamma}{1-\gamma}\frac{R(r)}{w(r)}(1-\tilde{\gamma}F(r)) \tag{B.24}$$

where $\tilde{\gamma} = \frac{\gamma/(1-\gamma)-\gamma_P/(1-\gamma_P)}{\gamma/(1-\gamma)}$. Land rents can be obtained from equation (B.24) as

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

whereas total income in city r is given by

$$\frac{1}{\gamma}w(r)n(r) = \frac{1}{\gamma}\frac{1-F(r)}{1-\tilde{\gamma}F(r)}w(r)N(r).$$

Using these results in the derivation of the equilibrium conditions, we obtain that the population of city r is the solution to the following equation:

$$\begin{aligned} N(r)^{[1+\eta\sigma+(1-\gamma-\alpha)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} =& \gamma^{\sigma-1}\tilde{a}(r)^{\frac{\sigma(\sigma-1)}{2\sigma-1}}A(r)^{\frac{(\sigma-1)^2}{2\sigma-1}}(1-\tilde{\gamma}F(r))^{[1+(1-\gamma)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} \cdot \\ & (1-F(r))^{-\frac{\sigma-1}{2\sigma-1}}MA(r) \end{aligned} \tag{B.25}$$

where

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

Equation (B.25) allows us to state the following proposition, which is the counterpart of Proposition 1 in Section 4.2.

Proposition 2. *If $\gamma_P < 1$, then an increase in the share of land allocated to the port in city in r, $F(r)$, decreases shipping costs $\mathbf{E}[T(r,s)]$, thus increasing $MA(r)$. Everything else fixed, an increase in $MA(r)$ increases the population of the city (market access effect). Holding $MA(r)$ fixed, if $\gamma_P \geq \gamma$, an increase in $F(r)$ draws additional people into the city (employment effect). If $0 < \gamma_P < \gamma$, an increase in $F(r)$ may trigger either a positive employment effect or migration out of the city (land use effect), depending on the values of structural parameters γ, γ_P and σ. If and only if $\gamma_P = 0$ (our baseline model), the model implies a negative land use effect irrespectively of the values of structural parameters. Finally, an increase in $F(r)$ increases shipping flows, thus lowering amenities $\tilde{a}(r)$ and, therefore, city population (disamenity effect).*

Proof. The results follow directly from equation (B.25). □

According to Proposition 2, an expansion of port activity has different implications on city population depending on labor's share in transshipment. Besides the standard market access and disamenity effects, port development affects city population in two ways. First, it draws people into the transshipment sector as long as labor's share in the sector is different from zero. Second, it decreases the amount of land available for the production of the city-specific good, which induces workers in this sector to leave the city. If labor's share in the transshipment sector is sufficiently high, the first effect always dominates the second one (employment effect). This implies that the population of the city should increase even more than what is implied by the standard market access effect. Such a positive effect, however, is not consistent with what we find in the data (Section 5), in particular, with the negative and significant coefficient on shipping once we control for market access. While a very strong disamenity effect could still, in principle, explain the negative impact of shipping on population, we find this effect to be quantitatively small (Section 6.3).

To sum up, the model presented in this section sheds light on two facts. First, if the share of labor in transshipment is too high, the model with transshipment labor has different implications than our baseline framework. These implications, however, are in contrast with the empirical findings of Section 5. Second, if the share of labor in transshipment is sufficiently low, the model with transshipment labor is more complex in its structure but delivers predictions that are similar to the predictions of our baseline framework.

B.7 A model with monopolistic competition in transshipment

This section presents a version of our baseline model in which landlords providing transshipment services engage in monopolistic competition. This implies that, unlike in our baseline model, port activity involves positive profits. We also show how we take the model with monopolistic competition to the data and how we simulate the same no-containerization counterfactual in it for one of the robustness exercises in Appendix B.8.

We first present the setup of the monopolistic competition model. As in our baseline model, we assume that each city is inhabited by a continuum of landlords. Without loss of generality, we normalize the mass of these landlords to one in each city, and index an individual landlord by $m \in [0,1]$.

Unlike in our baseline model, we assume that transshipment services are differentiated products. Firms shipping through port city r may use the services of any number of landlords m residing in the city. Firms aggregate transshipment services in a CES function with elasticity of substitution $\zeta \in (1, \infty)$ across the services performed for them by the individual landlords. As $\zeta < \infty$, these services are imperfect substitutes. Hence, each firm uses the transshipment service of each landlord in equilibrium.[10]

Landlords are aware that they are the sole provider of their differentiated transshipment service but cannot influence city-wide prices and quantities. Thus, they engage in monopolistic competition, choosing their land allocation, transshipment price and transshipment quantity to maximize their net nominal income. In other words, landlord m in port city r solves the problem

$$\max_{F_m(r), O_m(r), Shipping_m(r)} \left[O_m(r) - (\nu(r) + \psi(F_m(r)))\, Shipping(r)^{\lambda} \right] Shipping_m(r) + R(r)(1 - F_m(r))$$

where $O_m(r)$ is the price of transshipment services that landlord m charges, $\nu(r)$ is the exogenous part of transshipment costs, $F_m(r)$ is the share of land that the landlord allocates to transshipment, $Shipping(r)$ is the total value of shipments flowing through the port excluding the price of transshipment services, $R(r)$ is the land rent prevailing in the city, and $1 - F_m(r)$ is the share of land rented out to firms.

As the price elasticity of demand for each landlord's transshipment service is constant at $-\zeta$, each landlord charges a constant markup over her marginal cost in equilibrium:

$$O_m(r) = \frac{\zeta}{\zeta - 1} (\nu(r) + \psi(F_m(r)))\, Shipping(r)^{\lambda}$$

As landlords in a given port city are symmetric, we can drop their index and simply write

$$O(r) = \frac{\zeta}{\zeta - 1} (\nu(r) + \psi(F(r)))\, Shipping(r)^{\lambda} \tag{B.26}$$

from which we get that landlords earn profits on transshipment equal to

$$\Pi(r) = \frac{1}{\zeta - 1} (\nu(r) + \psi(F(r)))\, Shipping(r)^{1+\lambda}. \tag{B.27}$$

[10] To fix ideas, one may think that one port city landlord provides the cranes, another the storage, and so on. As a result, firms use the services of all landlords, not only one.

For simplicity, we assume that landlords spend these profits outside our set of cities S. This implies that we do not need to take profits into account when calculating demand for goods in the city, or city GDP. This assumption helps us keep the model computationally tractable.

The first-order condition to the landlord's maximization problem with respect to $F_m(r)$ implies

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

from which, by rearranging,

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

Note that this equation is identical to equation (B.5) of our baseline model. More generally, as the remaining model assumptions in the monopolistic competition model are the same as those in the baseline model, the only equation that differs between the two frameworks is equation (B.26), which replaces equation (6) in the baseline model. The remaining equilibrium conditions are all identical.

In Section B.8, we conduct a robustness check in which we take the model with monopolistic competition to the data to measure the aggregate and local effects of containerization. Inverting and simulating the monopolistic competition model follows the same steps as described in Appendix B.3 and Appendix B.4, with one exception: we use equation (B.26) instead of equation (6) whenever we calculate transshipment prices.

To do so, we need to choose the value of the markup parameter ζ. Note that, by equation (B.27), transshipment profits are decreasing in ζ. Data on profits of ports are hard to find, especially during our period of interest, but we were able to obtain profit and revenue data for a number of ports from annual reports of port authorities between 1950 and 1990.[11] In this sample, profits as a percentage of revenue are on average 28%, with no clear trends over time. Choosing $\zeta = 3$, our model predicts an average profit margin of 27% and a median profit margin of 33% across ports. Hence, we use $\zeta = 3$ in the inversion and the counterfactual simulation.[12]

B.8 Aggregate and local effects of containerization: robustness

In Table D.14, we examine the sensitivity of the model-implied aggregate and local effects of containerization to different values of the containerization shock and some alternative modeling choices. We focus on the sensitivity of our six headline findings: the aggregate welfare gains

[11] We describe these data in Section C.15.

[12] We compute the profit margin of port r in the model as $\frac{\Pi(r)-R(r)F(r)}{O(r)Shipping(r)}$. These margins vary across ports and are in fact negative for a few of them. As these ports operate in the data, we do not let them shut down in the model and assume they are subsidized from the outside economy.

from containerization; the aggregate resource costs; the aggregate specialization gains; the effect of depth on shipping (Stylized Fact 1); the reallocation of shipping toward less land-scarce cities (Stylized Fact 2); and the local population effects of shipping (Stylized Fact 3). Row (1) of Table D.14 repeats these results in our baseline model calibration, while rows (2) to (14) report them for each of our thirteen robustness exercises.

In the exercises of rows (2) and (3), we use higher and lower values of our transshipment cost parameter β, respectively. We argued in Section 6.1.2 that our endogenous port development mechanism is stronger under higher values of β. Thus, it is not surprising that we obtain higher resource costs and specialization gains from containerization, more reallocation of shipping toward less land-scarce cities, and more negative net population effects as a result of ports' more intensive land use in row (2). The opposite is true under the lower β of row (3).

In rows (4) and (5), we use alternative values of our counterfactual β: one that implies a smaller (65%) increase in the mean port share, and one that implies a larger (85%) increase. As expected, a smaller increase in land use leads to slightly higher welfare gains, lower resource costs and lower specialization gains from containerization.

In row (6), we do not offset the relationship between exogenous transshipment costs and port depth in the counterfactual. As depth no longer plays a role in the model in this case, we cannot estimate the stylized facts, as estimating these reduced-form coefficients relies on depth as an instrument for shipping changes. Nevertheless, we find aggregate welfare gains, resource costs and specialization gains that are fairly close to our baseline estimates.

In rows (7) and (8), we choose ν_{CF} to target different (30% and 20%, respectively) changes in the sum of exogenous and endogenous transshipment costs. Unsurprisingly, a larger change in total transshipment costs is associated with higher aggregate gains, resource costs and specialization gains. The opposite is true if we assume that total transshipment costs changed less. The estimated local effects of containerization hardly change, however.

To study how the assumption of perfect competition in transshipment influences our results, we develop a model in which the provision of transshipment services is subject to monopolistic competition in Appendix B.7. The key difference relative to our baseline setup is that transshipment activity involves positive profits in this monopolistic competition model. Row (9) reports the aggregate and local effects of containerization in the monopolistic competition model. While the results obviously change to some extent, they remain close to our baseline model, both qualitatively and quantitatively.

Rows (10) and (11) add a uniform 10% and 20% change in the elasticity of inland shipping costs to distance, respectively. This amounts to making inland shipping costs higher in the counterfactual, mimicking a decline in inland shipping costs brought about by containerization besides the change in transshipment costs. Obviously, adding an inland shipping cost reduction increases the estimated aggregate gains. However, the resource costs and specialization gains from containerization, as well as the local effects, remain very similar to our baseline model. The reason for this is that the overland cost reduction has two opposing effects of roughly similar magnitude. On the one hand, overland transport cost reductions make these routes more attractive relative to sea routes, leading to less endogenous port development. However, they also increase the overall volume of shipping, which increases port development.[13] Put differently, we find no evidence of missing interaction effects between the aspect of containerization we are interested in, and the overland transport cost reductions, which we do not account for in the main analysis of this paper.

Rows (12) to (14) study the robustness of our results to alternative versions of port city disamenities. In row (12), we assume that disamenities are present in 1990, but not in the pre-containerization equilibrium. In other words, we assume that it was containerization that brought about the disamenities associated with port activities. Unsurprisingly, loading all the negative disamenity effects on containerization substantially lowers the overall welfare gains from the new technology. At the same time, it leaves the estimated land use costs, specialization gains and local effects almost unchanged.

In row (13), disamenities are present both before and after containerization but the value of parameter ρ that disciplines their strength is twice as large as in the baseline calibration. As expected, stronger disamenities from port activities make the estimated local population effects more negative, so much so that the estimate becomes statistically significant, unlike in the data. By contrast, in row (14), we shut down disamenities completely, setting $\rho = 0$. Here, the estimated local population effects (Stylized fact 3) turn positive, though still insignificant. Overall, we conclude that the estimated aggregate and local effects of containerization are fairly stable across these different model specifications.

B.9 The declining labor intensity of port technology

Containerization is a labor-saving technology as the standardization of cargo-handling allows for more extensive automation (Bridgman, 2021). If the job destruction that occurred in port-related activities was sufficiently large, it could account for a substantial part of the negative effects on population that we identify in our empirics (Section 5). We assess the magnitude of this channel using readily available historical data on U.S. employment in water transportation (the industry to which longshoremen belong) from the *Bureau of Economic Analysis (BEA)*.[14]

There were approximately 222,000 full-time equivalent employees in water transportation in

[13]The canceling out of these two forces is also reflected in the fact that mean port size changes are very similar with and without overland cost changes; in both cases, the increase is about 8.6 percentage points.

[14]We use 'Water Transportation' employees from the BEA's 'Full-Time Equivalent Employees by Industry' series (Table 6.5). Accessed February, 2021 at https://apps.bea.gov/iTable/iTable.cfm?reqid=19&step=2&isuri=1&1921=survey.

the U.S., accounting for 0.12% of the population in 1960.[15] Between 1960 and 1987 (the last year in the series), employment fell by 23%. We conduct a back-of-the envelope calculation to assess the magnitude of population losses that this channel could account for. We make the following conservative assumptions: i) we assign *all* U.S. water-transportation workers pre-containerization to our sample of U.S. cities; ii) we assume *all* of these jobs were made redundant by containerization (when in fact, only 23% of jobs disappeared); and iii) we assume *all* workers moved out of the port city. This would result in a reduction of the population in port cities by 0.34% relative to 1960 population levels.[16]

These population losses are about an order of magnitude too small relative to i) the causally estimated average population loss, which is 3.8%,[17] or ii) the average land use effect implied by the quantitative model, which is 2.25%.[18] Both of these numbers are about an order of magnitude larger than the entire size of the water transport sector in 1960. Based on this, we conclude that the declining labor intensity of port technology can at most account for a very small fraction of the negative effects on population induced by containerization.

B.10 Elasticity of port city disamenities with respect to shipping

Welfare in Los Angeles, according to our model, is

$$u(LA) = GDP(LA) * \bar{a}(LA) * [1 + Shipping(LA)]^{-\rho}$$

Marquez and Vallianatos (2012) estimate the economic cost of pollution emitted by the ports of Los Angeles and Long Beach to be $30 billion annually (including economic cost associated with deaths as well as medical care for illnesses and missed school and work days). This is equivalent to 5% of the GDP of Los Angeles (averaged across 2002 and 2010, from *Canback*, see Appendix section C.4). Therefore,

$$1 - \frac{pollutioncost(LA)}{GDP(LA)} = (1 + Shipping(LA))^{-\rho}$$

Rearranging and plugging in the shipping for LA ports, which amounted to 7,959 vessels in 2005 for the ports of Los Angeles and Long Beach according to our own data, yields $\rho = 0.005$.

[15] Source: https://en.wikipedia.org/wiki/1960_United_States_census. Employees in water transportation accounted for 0.34% of all full-time equivalent employees.

[16] The population of the sample of U.S. Geopolis cities was 64,951,000 in 1960.

[17] Shipping increased on average by 24% between 1960 and 1990. Multiplying this by the elasticity of population w.r.t. shipping estimated in Table 3, column 2 (-0.159) yields a 3.8% population loss.

[18] In the quantitative estimation, the size of the non-shipping sector decreases on average by 8.98% across port cities with the introduction of containerization. The log-linearized equation determining equilibrium city population (equation 7) implies a population loss of 8.98%*0.25 = 2.25%, where the elasticity of 0.25 is based on the parameter values used in the quantitative estimation (and reported in Table D.9).

C Data

In this section, we provide additional details about data construction and sources for the variables used in the analysis.

C.1 Lloyd's List shipping data

We clean the shipping data by manually matching them to the 1953 and 2017 editions of the *World Port Index (WPI)*, which is a widely used reference list of worldwide ports. The initial Lloyd's List sample of 'ports' included ports on navigable rivers such as Budapest, Hungary. We therefore chose to discipline the sample of ports using WPI. We use a historic and current edition of the WPI to ensure we capture both ports that may no longer exist, and ones that only appear later in the period. A different approach would have been to choose a distance threshold from the coast and drop any port located further from the coast than the threshold. This definition, however, is very sensitive to the precision of the coastline shapefile used to calculate distance form the coast, which is why we did not choose this method. Despite filtering the Lloyd's List sample through the WPI, our final sample still contains a handful of ports that are very far inland. In the empirical analysis, we show that our results are robust to different ways of treating these 'inland ports.' Our base sample consists of Lloyd's List ports that match to at least one of the WPI editions.

C.2 Underwater elevation levels

We use data on underwater elevation levels from the *General Bathymetric Chart of the Oceans (GEBCO)*. We use the 2014 version of these data. Most observations in the dataset are from ship-track soundings with interpolation between soundings guided by satellite-derived gravity data. The data are continuously updated with sources from local bathymetry offices and coastal navigation charts. More details on dataset construction can be found at http://www.gebco.net.

C.3 Saiz land scarcity measure

The following sources are used to calculate the Saiz measure for our sample of cities. The coastline shapefile needed to distinguish between land and sea cells is from GSHHG (https://www.soest.hawaii.edu/pwessel/gshhg/). Inland bodies of water and wetlands are from the World Wildlife Fund's *Global Lakes and Wetlands Database*.[19] Finally, data on land elevations used to calculate the slope of each cell is from GEBCO's land data, described above.

C.4 Predicted city-level GDP per capita

Here we provide a more detailed discussion of how we estimate city level GDP per capita for our full sample of cities (port and non-port cities). First, 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*.

[19]Link: https://www.worldwildlife.org/pages/global-lakes-and-wetlands-database.

GDP are reported at purchasing power parity (in 2005 USD). We have estimates from this source for 898 cities in our sample.

We estimate city-level GDP for the full sample by extrapolating the estimated relationship between GDP per capita and nighlight luminosity. We begin by estimating the linear fit of GDP per capita on nightlight luminosity, building on a growing body of evidence suggesting that income can be reasonably approximated using nightlight luminosity data (Donaldson and Storeygard, 2016).

We construct the 'luminosity' of each city in the following way. We take the 1992 30 arc-second grid layer from NOAA's *National Geophysical Data Center* (source: https://ngdc.noaa.gov/products/) as the baseline input, as this is the closest year to 1990 – the year for which we have city income from *Canback*. We define a cell in this raster to be 'lit up' if its luminosity level is above 25. This threshold defines meaningful levels of economic activity in the cell - as proxied by nightlights.[20] We then construct a polygon from contiguous cells with luminosity above 25 for each city in our sample. We observe luminosity for 2,294 cities in our dataset.[21] With these data in hand, we then define a city i's luminosity, $luminosity_i$, to be the sum of all cells' luminosity levels within the polygon. Note that in this summation, we drop any cells identified as 'gas flares' in the source data, as these do not contain meaningful information on economic activity.

For the remaining 342 cities (13%), we either fail to identify an area polygon assigned to the city (340 cities) or a gas flare completely covers the polygon of the city (2 cities). We observe both GDP per capita and luminosity for a subset of 810 cities. For this subset, we estimate the relationship between GDP per capita and luminosity. More precisely, we estimate

$$\ln(GDP/capita)_i = \beta * \ln(luminosity_i) + FE_c + \epsilon_i \tag{C.28}$$

where $GDP/capita_i$ is city-level GDP per capita as compiled in the *Canback Global Income Distribution Database (CGIDD)* for the year 1990 which covers 898 cities, and $luminosity_i$ measures the sum of luminosity in the cells in the polygon that defines the area of the city.

Note that most 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 cities worldwide. To account for these differences and the way in which they affect luminosity, we include country fixed effects FE_c in our estimation. However, in order to identify country fixed effects, we need to drop 21 cities that are the only cities with GDP per capita data in their respective country, leaving a sample of 789 cities for the estimation.

The results of this regression are given in column (1) of Table D.17. We then predict GDP per

[20] We experimented with different cutoffs and this was the one for which the R^2 in the regression of income on luminosity was highest.

[21] We have cities with 'missing' luminosity data if we fail to detect *any* cells with luminosity levels above 25 in the vicinity of the city's geocode.

capita for all cities for which we observe luminosity that are also in the set of countries used in this regression. This allows us to predict GDP per capita for a total of 2,289 cities. For the remaining 341 cities, we use the following approximation. For 89 cities, we observe GDP per capita directly, which we use. For 240 cities, we only observe population in 1990, so we use this to predict GDP per capita based on the estimated relationship between GDP per capita and population in 1990 for all cities in our sample for which we observe both measures. This estimated relationship is given in column (2) of Table D.17. Finally, for 18 cities, we only observe population in 1980, so we use the latter to predict GDP per capita for all cities in our sample for which we observe both variables, resulting in the estimated relationship in column (3) of Table D.17.

This procedure yields a city-level estimate for GDP per capita for all 2,636 cities in our dataset.

C.5 Port shares for 1990

Here, we provide details on the construction of port share data and the sources used. First, it is important to note that historical data on the area occupied by the port are very difficult to find. For example, data on port area are only sporadically and inconsistently reported in *Lloyd's Ports of the World*, and they are usually not found in ports' annual reports. These are in fact the two sources from which we take the measure for the ports where port area is observed. We also experimented with using satellite images from the 1980s, but the resolution is too low to detect port areas.

We observe data on port area in 1990 for seven cities. These are: Aarhus (Denmark), Helsinki (Finland), Copenhagen (Denmark), Hamburg (Germany), Los Angeles (USA), New Orleans (USA) and Seattle (USA).[22] Data for the European ports and for the port of Los Angeles are from *Lloyd's Ports of the World* (1990). We complemented these with data for other U.S. ports where planning maps and annual reports gave information on the land area of the port. In all these cases, we verified or cleaned the data to ensure that a consistent definition of port area was used. In particular, these measures only include the total land (and not sea) area occupied by the port. Data for the remaining U.S. ports are from Port Authority of Seattle (1989) and Port of New Orleans (1984). These documents were shared by the port authorities based on requests we made. For Long Beach, we take port area in 1971 from the port's annual report (Port of Long Beach, 1971) and add additional land acquired from a detailed history of port projects (Riffenburgh, 2012). To construct the port shares, we use the area of *land* occupied by the city as reported in Wikipedia.

C.6 Area per throughput calculation for the Port of Seattle

We obtained 'Property Books' that allow us to calculate the area of the Port of Seattle from the *Port of Seattle Public Records Office*. These volumes contain engineering maps for each parcel of land under the ownership of the port. Each map includes an estimate of the land area. For both years 1961 and 1973, we used only land parcels directly related to port activities. In particular,

[22] The port area for Los Angeles includes the area occupied by the ports of Los Angeles and Long Beach.

we excluded the airport and the marina terminal. Data on annual total throughput (in short tons, including both domestic and international sea-borne trade) and the share of containerized cargo were collected from *Annual Reports* that are archived at the *Puget Sound Regional Archives*. To smooth out fluctuations in year-to-year capacity utilization, we took the five-year moving-average of throughput.

Table D.18 reports the numbers. While the expansion of traffic during this period was impressive (throughput doubled), the area occupied by the port expanded even more rapidly (increasing almost fourfold), such that area per throughout increased by 90% during this period. The *Annual Reports* paint a consistent picture. In the early 1960s, the port acquired vast parcels of land in the Lower Duwamish Industrial Development District. Throughout the latter half of the decade, the port continued to acquire more land in this area and to simultaneously develop the acquired tracts. These were completed in the late 1960s, early 1970s. We illustrate this in Figure E.8 which shows the set of acquired land parcels and an example of a completed container facility.

C.7 Google Earth port area and containerization, modern data

We compiled data on the area of all ports for a random subset of port cities in our dataset (236 cities, which is 43% of the full sample), resulting in 252 individual ports. For each port, we hand-coded polygons that contain port activities based on satellite imagines from *Google Earth*. We used the name tags of buildings as well as visual markers (e.g., stacked containers, ships). We aimed to be conservative in that we only included areas that could clearly be identified as containing port-related activities. As such, we did not include warehouses (as they cannot be unambiguously identified) or highways or railways. A port can have multiple polygons, e.g., in the case of terminals that are not directly connected. *Google Earth* reports the area (in km^2) of each polygon, which we aggregate to the level of Geopolis port cities. The average area of a port in our data is 3.6 km^2 (median: 2 km^2), with a minimum of 0.03 km^2 and a maximum of 30 km^2 (Los Angeles, including the Port of Long Beach). The latter occupies 43 km^2 according to Wikipedia (https://en.wikipedia.org/wiki/Port_of_Long_Beach and https://en.wikipedia.org/wiki/Port_of_Los_Angeles), so while our measure most likely underestimates the true size of ports, the measure is arguably in the correct range.

Data on total (in tons) and containerized (in TEUs; twenty-foot equivalent units) volume of cargo handled by each port is taken from the 2009 edition of *Le Journal de la Marine Marchande (JMM)*. We use the average of the reported numbers for 2008 and 2009 in order to maximize the number of observations, as some ports only report data for one of the two years. In order to generate the share of container traffic in total merchandise traffic, we use the average weight per TEU of 12 tons as recommended by the *European Sustainable Shipping Forum*.[23]

[23] Downloaded on March 11, 2021, from https://ec.europa.eu/clima/sites/clima/files/docs/0108/20170517_guidance_cargo_en.pdf.

We match the dataset on the area of ports and cargo volume based on the names, countries and geocodes of the ports, resulting in 123 observations.

C.8 Land reclamation

Data on land reclaimed from the sea are taken from Martín-Antón et al. (2016). The authors compare historical maps to current Google Earth images to examine whether land reclamation has taken place in a city. We matched these data to our sample of port cities. The authors report three measures; i) any land reclamation, ii) coastal land reclamation, iii) coastal and island land reclamation. This contains land reclaimed for any purpose, not just for port activities. In our analysis, we use their coastal land reclamation measure, though the results are essentially the same regardless of the measure used.

The authors systematically examined the coastlines of the world, paying particular attention to South East Asia, the Persian Gulf, Europe and the U.S., where land reclamation has been more extensive. Any systematic measurement error introduced in this way will be accounted for in our specifications that control for continent and coastline fixed effects. Reassuringly, the coefficients of interest do not change substantially with the inclusion of these, suggesting that these issues – if present – are not quantitatively large.

C.9 Country GDP per capita

Data on country-level GDP per capita are from the *Penn World Tables*. We take real GDP at constant 2011 prices (USD) and divide by country population reported from the same source. In theory, the data exist for 1950 (our first sample year), but in practice there are many missing observations. For this reason, in robustness checks, we always use the data for 1960. This is observed for many, though not all, countries.

C.10 Identifying city centroids for within port-city moves

In Section 3, we discussed evidence that showed that ports had moved further towards the outskirts of the city during our sample period. To conduct this exercise, we use data on ports' geocodes from two editions of the *World Port Index*: 1953 and 2017. We also need to identify the geocode of each city's centroid. To this end, we use daylight satellite data to identify a city's contiguous built-up area and find the city centroid within this polygon. We closely follow the methodology in Baragwanath, Goldblatt, Hanson, and Khandelwal (2019). In particular, we use an extremely high resolution dataset of daylight satellite data, the *Global Human Settlement Built-Up Grid* available at 38 m resolution (source: https://ghsl.jrc.ec.europa.eu/ghs_bu.php). Using this raster and the geocodes of our cities, we construct a polygon for each city consisting of contiguous built-up cells around the geocode. We take the centroid of this polygon to be the centroid of the city.

C.11 Ship size data

The evolution of ship sizes, illustrated in Figure E.9, is based on data purchased from the *Miramar Ship Index* (Haworth, 2020), accessible at http://www.miramarshipindex.nz. The *Miramar Ship Index* is a comprehensive list of all newly built ships and their main characteristics going back to the 19th century. We calculate the average tonnage of all newly built ships in the years 1960, 1990, and 2010, distinguishing between container-ships and non-container ships.

C.12 Nautical maps for dredging dummy variable

We obtained access to nautical maps of ports around the world from *marinetraffic.com*, see https://www.marinetraffic.com/en/online-services/single-services/nautical-charts.

These detailed nautical charts have been constructed based on information from hydrographic organizations of different countries. They provide pilotage information including depth of water at high spatial disaggregation. Dredged channels are demarcated on these maps by a 'safety contour' that distinguishes the channel from the surrounding shallow waters (defined as less than 5 meters). We constructed a binary variable, *'Dredging'*, that takes the value 1 if a dredged channel is visible on the nautical chart in the 3-5 km buffer ring around the port.

C.13 Port cost data based on Blonigen and Wilson (2008)

Blonigen and Wilson (2008) estimate port costs 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 U.S. census data for 1991 (see Blonigen and Wilson (2008) for additional details). The exporter fixed effects are all estimated relative to the port costs at Rotterdam. For our purposes, these relative measures need to be scaled to levels. We do this by setting the iceberg trade cost of passing through Rotterdam to be 1.004. This is based on estimates of revenue from handling one container to be approximately $140 AUD (Australian Competition and Consumer Commission, 2017, p. 8) and the average value of a container to be 20,000 EUR (Kirchner, 2006, p. 4).[24]

C.14 Data on frost-free days

We use data from the FAO GAEZ database (http://www.fao.org/nr/gaez/en/) to measure the average the number of frost-free days per year in each city. This database provides the average of this variable during the years between 1961 and 1990 in every cell over a 5 by 5 arc minute grid of the Earth. Using the geocoordinates of each city, we determine the grid cell in which the city is located, and assign the average number of frost-free days in the cell to the city.

[24]These are industry-level averages as of 2016 (for revenue from container handling) and 2006 (for average value of cargo), and do not refer specifically to data from Rotterdam.

C.15 Annual reports for ports

We were able to acquire annual reports for a number of port authorities in the United States during our sample period, 1950 to 1990, and for a handful of ports worldwide. Some ports have made historical annual reports available online, while for others, we have obtained the reports by contacting the port authorities. We use these reports i) for historical evidence (Section 1), ii) in the case of the Port of Seattle, to measure changes in land per unit of throughput during the period in which they containerized (Section 3.2), and iii) to calculate profit rates (Section B.7).

As accounting and reporting standards changed across ports and over time, we only kept ports that reported consistent information on profits over time (defined as revenue minus operating expenses and depreciation). These ports are: Houston, Los Angeles, Long Beach, New York/New Jersey, New Orleans, Seattle and Townsville (Australia). We tried to collect at least one observation per port for each decade between 1950 and 1990, and ended up with on average three decadal observations per port. The average profit margin across all observations in our sample is 28%, with no clear time trend. Data sources are as follows;

Houston. Port of Houston Authority of Harris County, Texas: 'Comprehensive Annual Financial Report' (various years). Thank you to Dollores Villareal at the Port of Houston for responding to our request and digitizing the data for us.

Los Angeles. Port of Los Angeles Board of Harbor Commissioners: 'Annual Report' (various years). Thank you to Kurt Arendt at the Port of Los Angeles for responding to our request and sharing data.

Long Beach. The Port of Long Beach California: 'Harbor Highlights' (various years). Accessible at https://www.polb.com/port-info/history#historical-publications.

New York/New Jersey. The Port Authority of New York and New Jersey: 'Annual Report' (various years). These can be accessed online at https://corpinfo.panynj.gov/pages/annual-reports/.

New Orleans. Board of Commissioners of the Port of New Orleans: 'Annual Report Fiscal' (various years). Thank you to Mandi Venderame at the Port of New Orleans for responding to our request and sharing data.

San Francisco. The Port of San Francisco: 'Annual Report', other reports and planning maps from various years. Thank you to Randolph Quezada at the Port of San Francisco for numerous helpful conversations and for sharing scans.

Seattle. The Port of Seattle: 'Annual Report' (various years) and planning maps. Thank you to Midori Okazaki, archivist at Puget Sound Regional Archives, for scanning the files during the COVID-19 lockdown while the archives were closed to the public.

Townsville (Australia). Townsville Harbor Board: 'Report' (various years). Thank you to the Port

Authority for responding to our data request.

D Tables

Table D.1: Summary statistics

	Observations	Mean	Standard Deviation
Shipment (annualized)	2,765	2,913	7,051
Population (in ’000s): *All Cities*	12,698	386	1,086
Population (in ’000s): *Port Cities*	2,735	724	1,886
Depth	553	2.19	1.49
Saiz land scarcity measure	553	0.44	0.19

Notes: Shipment reports the annualized flow of shipments across all port city – year pairs (in levels). Population refers to the level of the population of each city-year pair in thousands. Depth and the Saiz land scarcity measure are time invariant measures and are defined in the main text.

Table D.2: Relationship between dredging and measured depth

	Dredging		
Independent variables	(1)	(2)	(3)
Depth	-0.058**	-0.042*	-0.028
	(0.025)	(0.024)	(0.028)
Observations	100	100	100
R-squared	0.059	0.138	0.250
FE	none	continent	coastline

Notes: This table tests the extent to which the baseline measure of depth captures naturally endowed depth (as opposed to depth attained by dredging). Dredging is a binary indicator that takes the value of one if nautical maps from *marinetraffic.com* show the presence of a dredged channel. Depth is the baseline measure of port suitability used in the paper. The sample consists of 100 randomly selected ports from the baseline sample. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table D.3: Robustness to data choices

Panel A: Depth predicts shipping flows (Stylized fact 1)						
	ln(Shipment)					
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Shipment +1	IHST	Port Cities	Depth=0	100K
Depth × post 1970	0.247***	0.144***	0.164***	0.218***	0.247***	0.285***
	(0.059)	(0.029)	(0.034)	(0.060)	(0.059)	(0.071)
Observations	2765	2765	2765	2640	2765	1565
R-squared	0.126	0.156	0.155	0.133	0.126	0.139
Number of cities	553	553	553	528	553	313

Panel B: Containerization increased shipping more in less land-scarce cities (Stylized fact 2)						
	ln(Shipment)					
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Shipment +1	IHST	Port Cities	Depth=0	100k
Depth × post 1970	0.566***	0.318***	0.368***	0.583***	0.566***	0.348**
	(0.152)	(0.079)	(0.090)	(0.147)	(0.152)	(0.177)
Depth × Saiz × post 1970	-0.707**	-0.408**	-0.472***	-0.779**	-0.707**	-0.203
	(0.323)	(0.159)	(0.183)	(0.315)	(0.323)	(0.376)
Saiz × post 1970	0.975	0.740**	0.814*	0.950	0.975	0.694
	(0.804)	(0.376)	(0.436)	(0.802)	(0.804)	(0.963)
Observations	2765	2765	2765	2640	2765	1565
R-squared	0.129	0.161	0.159	0.137	0.129	0.139
Number of cities	553	553	553	528	553	313

Panel C: The local causal effect of shipping on population (Stylized fact 3)						
	ln(Population)					
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Shipment +1	IHST	Port Cities	Depth=0	100K
ln(Shipment)	0.015	0.027	0.024	0.025	0.015	0.045
	(0.049)	(0.086)	(0.076)	(0.053)	(0.049)	(0.052)
Observations	2734	2734	2734	2609	2734	1563
R-squared	0.717	0.719	0.719	0.720	0.717	0.606
Number of cities	552	552	552	527	552	313

Notes: 'Baseline' reports the baseline specification for comparability. Columns (2)-(3) examine robustness to different ways of dealing with zero shipping flows. Column (2) uses $\ln(Shipment + 1)$ as dependent variable – that is, we take the raw shipping variable and replace the zeros with ones and then take the natural logarithm. Column (3) uses the inverse hyperbolic sine transformation (IHST) for shipment. Different to the baseline, neither of these transformations annualizes the data. Columns (4)-(5) examine robustness to different ways of dealing with 'inland ports.' Column (4) drops them, reducing the sample size. Column (5) assigns depth equal to zero to these cities. Column (6) uses the subset of cities that already attained 100,000 inhabitants in 1950 to examine the effect of sample selection bias. 'Depth' indicates the port suitability measure interacted with indicators for decades including and after 1970. Standard errors clustered at the city level. *** p<0.01, ** p<0.05, * p<0.1.

Table D.4: Relationship between containerization and port area

	Ln(Port area, km2)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ln(Container traffic, TEUs)	0.288***	0.127***	0.133***	0.151***	0.153***	0.144***	
	(0.049)	(0.045)	(0.044)	(0.047)	(0.046)	(0.053)	
Ln(Total merchandise traffic, tons)		0.375***	0.283*	0.311***	0.247	0.356***	0.506***
		(0.080)	(0.166)	(0.080)	(0.161)	(0.118)	(0.069)
Ln(Non-bulk traffic, tons)			0.014		0.008		
			(0.099)		(0.096)		
Ln(Country GDP/capita)				0.311***	0.292**		
				(0.108)	(0.134)		
Container traffic share							0.562***
							(0.209)
Observations	123	123	73	122	73	123	123
R-squared	0.287	0.395	0.327	0.431	0.352	0.672	0.398
% change							0.75
Country FEs	×	×	×	×	×	✓	×

Notes: 'Container traffic share' is defined as $\frac{\text{container traffic in TEU-s}*12\text{ tons per TEU}}{\text{total merchandise traffic in tons}}$. Non-bulk traffic is all traffic net of liquid and solid bulk. Container traffic and total merchandise traffic are averaged across 2008 and 2009 in order to maximize the sample size. Country level GDP per capita and non-bulk traffic are for 2009. Data sources: *Google Earth* and *Le Journal de la Marine Marchande*. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table D.5: Shipping increased more in less land-scarce cities (Stylized fact 2)

	ln(Ship)			
Independent variables	(1)	(2)	(3)	(4)
Depth × post 1970	0.464***	0.566***	0.437***	0.497***
	(0.138)	(0.152)	(0.142)	(0.166)
Depth × Saiz land scarcity × post 1970	-0.408*	-0.707**	-0.431	-0.586*
	*-0.122**	*-0.211***		
	(0.220)	(0.323)	(0.308)	(0.331)
	{0.153}	{0.237}		
Saiz land scarcity × post 1970		0.975	-0.052	1.176
		(0.804)	(0.811)	(0.749)
Observations	2765	2765	2765	2360
R-squared	0.128	0.129	0.250	0.143
Number of cities	553	553	553	472
Year FE	✓	✓	✓	✓
City FE	✓	✓	✓	✓
Population 1950 × Year	✓	✓	✓	✓
Coastline × Year FE	×	×	✓	×
GDP pc (country) × Year	×	×	×	✓

Notes: 'Depth' indicates the port suitability measure. 'Saiz land scarcity' is the Saiz land scarcity measure defined in Saiz (2010). Each measure is interacted with an indicator for decades including and after 1970. Standardized coefficients in italics underneath the baseline coefficients. Standard errors clustered at the city level in parentheses. Conley standard errors to adjust for spatial correlation in curly brackets. *** p<0.01, ** p<0.05, * p<0.1 (significance refers to clustered standard errors).

Table D.6: Relationship between coastal land reclamation and the Saiz land scarcity measure

	Coastal land reclamation (indicator)					
Independent variables	(1)	(2)	(3)	(4)	(5)	(6)
Saiz land scarcity measure	0.1296*	0.1356**	0.1146			
	(0.0686)	(0.0678)	(0.0754)			
Depth				0.0008	0.0038	-0.0003
				(0.0093)	(0.0096)	(0.0106)
Observations	553	553	553	553	553	553
R-squared	0.00534	0.08521	0.13287	0.00001	0.07991	0.12925
FE	none	continent	coastline	none	continent	coastline

Notes: Dependent variable is equal to one if coastal land reclamation was reported, and zero otherwise. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table D.7: The local causal effect of shipping on population – full panel specification

	Panel Regression				
	ln(Pop)	ln(Ship)	ln(Pop)	ln(Ship)	ln(Pop)
Independent variables	(1)	(2)	(3)	(4)	(5)
ln(Ship)	0.015				
	0.035				
	(0.049)				
	{0.039}				
Depth × post 1970		0.268***	0.004		
		*0.143****	*0.005*		
		(0.058)	(0.013)		
Depth × 1960				-0.042	-0.003
				(0.064)	(0.008)
Depth × 1970				0.246***	0.007
				(0.069)	(0.013)
Depth × 1980				0.213***	-0.002
				(0.079)	(0.017)
Depth × 1990				0.280***	0.002
				(0.086)	(0.020)
Observations	2734	2734	2734	2734	2734
Number of cities	552	552	552	552	552
Year FE	✓	✓	✓	✓	✓
City FE	✓	✓	✓	✓	✓
Population 1950 × Year	✓	✓	✓	✓	✓
Population 1950	×	×	×	×	×
Specification	2SLS	FS	RF	dyn FS	dyn RF
KP F-stat	21.13				

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

Table D.8: The local causal effect of shipping on population – robustness

	ln(Population)			
Independent variables	(1)	(2)	(3)	(4)
ln(Ship)	0.015	-0.071	0.018	-0.015
	(0.049)	(0.060)	(0.051)	(0.051)
Observations	2734	2734	2734	2338
R-squared	0.717	0.759	0.717	0.756
Number of cities	552	552	552	471
Year FE	✓	✓	✓	✓
City FE	✓	✓	✓	✓
Population 1950 × Year	✓	✓	✓	✓
Coastline × Year FE	×	✓	×	×
Saiz × Year	×	×	✓	×
GDP pc (country) × Year	×	×	×	✓
Specification	2SLS	2SLS	2SLS	2SLS
KP F-stat	21.13	13.71	16.26	19.48

Notes: All specifications are 2SLS, using the depth measure as an instrument for shipping (interacted with a dummy for decades including and after 1970). Standard errors clustered at the city level. *** p<0.01, ** p<0.05, * p<0.1.

Table D.9: Calibration of structural parameters

Parameter	Target
$\alpha = 0.06$	Agglomeration externalities (Ciccone and Hall, 1993)
$\gamma = 0.84$	Non-land share in production (Desmet and Rappaport, 2017)
$\eta = 0.15$	Migration elasticity (Kennan and Walker, 2011)
$\sigma = 4$	Elasticity of substitution across tradables (Bernard et al., 2003)
$\rho = 0.005$	Elasticity of disamenities with respect to shipping (see section B.10)
$\theta = 203$	Idiosyncratic shipping cost dispersion (Allen and Arkolakis, 2019)
$\lambda = 0.074$	Congestion externalities in ports (Abe and Wilson, 2009)

Table D.10: Prediction of population based on the number of frost free days

	ln(Population)								
			Africa	North America	Latin America	Asia	Europe	Oceania	USSR
Independent variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Frost Free Days × 1960	0.0007***	0.0003*	-0.0005	0.0015***	0.0008***	-0.0001	-0.0001	0.0014	0.0001
	(0.0001)	(0.0002)	(0.0013)	(0.0003)	(0.0002)	(0.0002)	(0.0001)	(0.0012)	(0.0007)
Frost Free Days × 1970	0.0017***	0.0006***	-0.0006	0.0026***	0.0013***	0.0008***	0.0004**	0.0023	0.0012
	(0.0001)	(0.0003)	(0.0021)	(0.0005)	(0.0003)	(0.0002)	(0.0002)	(0.0020)	(0.0010)
Frost Free Days × 1980	0.0028***	0.0012***	-0.0006	0.0039***	0.0017***	0.0005*	0.0011***	0.0036	0.0010
	(0.0001)	(0.0003)	(0.0020)	(0.0006)	(0.0004)	(0.0003)	(0.0003)	(0.0027)	(0.0011)
Frost Free Days × 1990	0.0039***	0.0013***	-0.0009	0.0047***	0.0015***	0.0014***	0.0013***	0.0049	0.0011
	(0.0002)	(0.0004)	(0.0026)	(0.0007)	(0.0005)	(0.0003)	(0.0003)	(0.0031)	(0.0012)
Observations	12368	12368	1184	987	1532	4784	2410	104	1367
R-squared	0.729	0.839	0.798	0.686	0.852	0.756	0.556	0.647	0.762
Number of cities	2568	2568	245	198	308	1038	482	21	276
City FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country year FE	×	✓	×	×	×	×	×	×	×
Year FE	✓	×	✓	✓	✓	✓	✓	✓	✓

Notes: Column (2), which includes country-year fixed effects, is our preferred specification used for predicting population. Column (1) only controls for year fixed effects. The time pattern in the effect of frost free days on population is similar; however, country-year fixed effects take out some of the explanatory power of temperature and are therefore a more conservative measure. Columns (3) to (9) estimate column (1) for countries in different regions of the world. Reassuringly, the number of frost free days has the strongest effect in North America, where air conditioning is arguably most prevalent, has medium effects in Latin America, Asia, and Europe, and has no detectable effects in Africa, Oceania and USSR, where air conditioning was probably less wide-spread. Frost Free Days indicates the average number of frost free days per year in the city between 1961-1990. Standard errors clustered at the city level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table D.11: Model-inspired specification: fully flexible first stages

Independent variables	(1) ln(Ship)	(2) ln(Market Access)
Depth × 1960	-0.038	0.003
	(0.064)	(0.002)
Depth × 1970	0.251***	0.007***
	(0.069)	(0.002)
Depth × 1980	0.226***	0.009***
	(0.079)	(0.002)
Depth × 1990	0.290***	0.009***
	(0.085)	(0.002)
Market Access IV	7.169	1.929***
	(5.413)	(0.140)
Observations	2696	2696
R-squared	0.125	0.728
Number of cities	544	544
Year FE	✓	✓
City FE	✓	✓
Population 1950 × Year	✓	✓

Notes: 'Depth' indicates the port suitability measure. It is interacted with dummy variables for decades in order to examine the time path of when depth started to matter for $\ln(Ship)$ and $\ln(MarketAccess)$. '$\ln(MarketAccess)$' is the empirical counterpart of the market access term, defined in Section 5. 'Market access IV' is the instrument for the market access term, defined in Section 5. Standard errors clustered at the city level. *** p<0.01, ** p<0.05, * p<0.1.

Table D.12: Model-inspired specification – robustness

	ln(Population)							
Independent variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(Ship)	-0.159**	-0.156**	-0.147**	-0.084**	-0.080	-0.164**	-0.231***	-0.215***
	(0.065)	(0.067)	(0.072)	(0.041)	(0.058)	(0.074)	(0.066)	(0.081)
ln(Market Access)	7.103***	6.982***	6.613***	0.588	7.111***	5.692***	10.090***	9.400***
	(0.795)	(0.844)	(1.043)	(2.918)	(0.713)	(1.354)	(1.250)	(1.514)
Observations	2696	2696	2696	2696	2696	2303	2696	2696
R-squared	0.417	0.429	0.467	0.755	0.507	0.544	(1.250)	(1.514)
Number of cities	544	544	544	544	544	464	544	544
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
City FE	✓	✓	✓	✓	✓	✓	✓	✓
Population 1950 × Year	✓	✓	✓	✓	✓	✓	×	×
Coastline × Year FE	×	×	×	✓	×	×	×	×
Saiz × Year	×	×	×	×	✓	×	×	×
GDP pc (country) 1960 × Year	×	×	×	×	×	✓	×	×
Specification	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Drop Cities in Market Access IV	none	≤ 200	≤ 500	none	none	none	none	none
Borusyak Hull correction	none	none	none	none	none	none	worldwide	30° lat
KP F-stat	9.63	9.16	7.75	4.02	8.64	8.43	12.98	8.694

Notes: This table reports the same specification as Table 3, column (2). All specifications are 2SLS. The instruments used are the depth measure and the market access IV as defined in Section 5. Nearby cities are dropped from the market access IV in columns (2)-(3). 'Drop Cities in Market Access IV' defines the point-to-point distance buffer (in km) for the set of cities to be dropped. Columns (4)-(6) examine robustness of the results to the inclusion of the same set of controls we added in Section 3. Columns (7)-(8) use Borusyak and Hull (2022) to correct for potential non-random shock exposure in the market access IV by reshuffling the number of frost free days across cities. 'Borusyak Hull correction' specifies how the reshuffling is done. Standard errors clustered at the city level. *** p<0.01, ** p<0.05, * p<0.1.

Table D.13: Heterogeneity of the effect of land scarcity on shipping – model-simulated data

Independent Variables	(1) Δ ln(Ship)
Depth	0.513***
	(0.084)
ln(Rent$_{CF}$) × Depth	-0.022**
	(0.009)
ln(Rent$_{CF}$)	0.013
	(0.012)
Observations	553
R-squared	0.528

Notes: 'Depth' indicates the port suitability measure. $\ln(Rent_{CF})$ refers to the logarithm of the counterfactual (pre-containerization) land rents. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table D.14: The aggregate welfare effects of containerization – sensitivity analysis

Model	Welfare effect (%)	Resource cost (pp)	Specialization gains (pp)	Stylized fact 1	Stylized fact 2	Stylized fact 3
1. Baseline	3.23	-0.64	0.39	0.262*** (0.011)	-0.013** (0.005)	-0.030 (0.048)
2. 20% higher β in inversion	3.20	-0.77	0.45	0.267*** (0.012)	-0.014** (0.006)	-0.050 (0.049)
3. 20% lower β in inversion	3.31	-0.51	0.33	0.258*** (0.010)	-0.012** (0.005)	-0.011 (0.048)
4. Counterfactual β implying 65% increase in port share	3.24	-0.61	0.36	0.263*** (0.011)	-0.013** (0.005)	-0.041 (0.048)
5. Counterfactual β implying 85% increase in port share	3.22	-0.67	0.41	0.261*** (0.011)	-0.013** (0.005)	-0.021 (0.048)
6. No depth-dependent change in $\nu(r)$	3.70	-0.79	0.50	n/a	n/a	n/a
7. Larger ν_{CF}: implies 30% change in total transshipment cost	3.96	-0.68	0.42	0.258*** (0.011)	-0.012** (0.006)	-0.011 (0.051)
8. Larger ν_{CF}: implies 20% change in total transshipment cost	2.50	-0.60	0.37	0.265*** (0.011)	-0.014*** (0.005)	-0.057 (0.046)
9. Monopolistic competition	3.43	-0.67	0.45	0.247*** (0.012)	-0.014** (0.006)	-0.008 (0.054)
10. Additional 10% inland cost transport reduction	4.77	-0.65	0.40	0.259*** (0.011)	-0.014*** (0.005)	-0.004 (0.048)
11. Additional 20% inland cost transport reduction	6.24	-0.65	0.40	0.256*** (0.011)	-0.015*** (0.005)	0.020 (0.050)
12. Disamenities only in inversion	1.37	-0.66	0.38	0.261*** (0.011)	-0.013** (0.005)	0.057 (0.050)
13. Disamenities twice as large	3.13	-0.64	0.41	0.262*** (0.011)	-0.013** (0.005)	-0.114** (0.047)
14. No disamenities	3.34	-0.65	0.39	0.262*** (0.011)	-0.013** (0.005)	0.059 (0.049)

Notes: Welfare effect refers to the gain in welfare due to containerization in our baseline model. Resource cost refers to the difference in welfare gains between Benchmark 1 (with exogenous and free transshipment cost reductions) and Benchmark 2 (with identical land use across port cities). Specialization gains refer to the difference in welfare gains between the baseline and Benchmark 2. The specifications for the remaining columns are the same as in column (2) of Table 4. Notice that the stylized facts cannot be estimated in model 6, as depth is not used in that counterfactual. *** p<0.01, ** p<0.05, * p<0.1.

Table D.15: The effects of targeted port development: The Maritime Silk Road

	Baseline				Benchmark 1			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Δ ln(Ship)	Δ ln(Port cost)	Δ ln(Market access)	Δ ln(Population)	Δ ln(Ship)	Δ ln(Port cost)	Δ ln(Market access)	Δ ln(Population)
Treated port city	0.78637***	-0.14993***	0.02936***	-0.02210***	0.60754***	-0.10536	0.02722***	0.00505**
	(0.12006)	(0.03158)	(0.00558)	(0.00598)	(0.09078)	0	(0.00583)	(0.00241)
Untreated port city in treated country	-0.41470***	0.01513**	0.01348***	0.01097***	-0.29817***	0	0.01526***	-0.00234
	(0.08222)	(0.00705)	(0.00285)	(0.00415)	(0.04999)	0	(0.00283)	(0.00218)
Port city in untreated country	0.01047	-0.00209	0.00076**	-0.00132	0.00292***	0	0.00084***	0.00021***
	(0.01109)	(0.00799)	(0.00031)	(0.00097)	(0.00034)	0	(0.00004)	(0.00003)
Inland city in treated country			0.02072***	0.00227***			0.02019***	-0.00013
			(0.00123)	(0.00066)			(0.00122)	(0.00023)
Inland city in untreated country			0.00041***	0.00004			0.00041***	-0.00003***
			(0.00006)	(0.00003)			(0.00005)	(0.00001)
Observations	553	544	2636	2636	553	553	2636	2636
R-squared	0.192	0.029	0.444	0.033	0.430	1.000	0.465	0.025

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

Table D.16: Maritime Silk Road Counterfactual – country fixed effects

	Baseline				Benchmark 1			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Δ ln(Ship)	Δ ln(Port cost)	Δ ln(Market access)	Δ ln(Population)	Δ ln(Ship)	Δ ln(Port cost)	Δ ln(Market access)	Δ ln(Population)
Treated port city	1.09616***	-0.13066***	0.01597***	-0.02959***	0.90913***	-0.10536***	0.01281***	0.00860***
	(0.10480)	(0.01257)	(0.00280)	(0.00879)	(0.06376)	0	(0.00234)	(0.00308)
Untreated port city in treated country	-0.92762***	0.00936	0.04693***	0.01060	-0.86417***	0	0.05048***	-0.00628
	(0.03864)	(0.00861)	(0.00596)	(0.01040)	(0.07243)	0	(0.00631)	(0.00820)
Port city in untreated country	0.00364***	-0.00006*	-0.00138***	-0.00154	0.00262***	0	-0.00132***	0.00029***
	(0.00045)	(0.00003)	(0.00031)	(0.00121)	(0.00021)	0	(0.00005)	(0.00007)
Inland city in treated country			0.05334***	0.00054			0.05507***	-0.00375
			(0.00482)	(0.00819)			(0.00516)	(0.00668)
Inland city in untreated country			-0.00191***	0.00009			-0.00154***	0.00006
			(0.00005)	(0.00006)			(0.00005)	(0.00006)
Observations	553	544	2636	2636	553	553	2636	2636
R-squared	0.639	0.378	0.942	0.186	0.923	1.000	0.979	0.210

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

Table D.17: Relationship between GDP per capita and nightlight luminosity

	ln(GDP per capita)		
Independent variables	(1)	(2)	(3)
ln(Luminosity)	0.126***		
	(0.014)		
ln(Population, 1990)		0.107***	
		(0.013)	
ln(Population, 1980)			0.100***
			(0.014)
Observations	789	854	871
R-squared	0.926	0.923	0.921
Country FE	✓	✓	✓

Notes: All regressions include country fixed effects. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table D.18: Port of Seattle: area per unit of cargo shipped

Year	Area	Throughput	Area/Throughput
1961	8,651,016	2,022,192	4.28
1973	33,547,908	4,135,795	8.11

Notes: Area reported in square feet, throughput in short tons. Data were not available far enough back in time to allow for the calculation of the five-year moving-average for 1961.

Table D.19: Prediction of port cost

	(1)
Independent Variables	Port cost
ln(Ship)	-0.033**
	(0.015)
Constant	0.444***
	(0.145)
Observations	72
R-squared	0.074

Notes: The dependent variable, port cost, is taken from the port efficiency estimation in Blonigen and Wilson (2008) for 1991, available for 72 international port cities in our data (for details, see Appendix C.13). The regressor, $\ln(Ship)$, refers to our shipping data in 1990. Observations are weighted by the inverse of the squared standard error of the estimated port cost as given by Blonigen and Wilson (2008). Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table D.20: Predicted port cost

	N	mean	standard deviation
1950	2145	0.345	0.105
1960	2145	0.328	0.106
1970	2145	0.324	0.108
1980	2145	0.317	0.105
1990	2145	0.294	0.104

Notes: This table shows summary statistics for the predicted port cost based on the estimation in Table D.19. The 2,145 ports include the 553 ports with population data (Geopolis ports) as well as all other ports from the Lloyd's List data that do not have population data.

E Figures

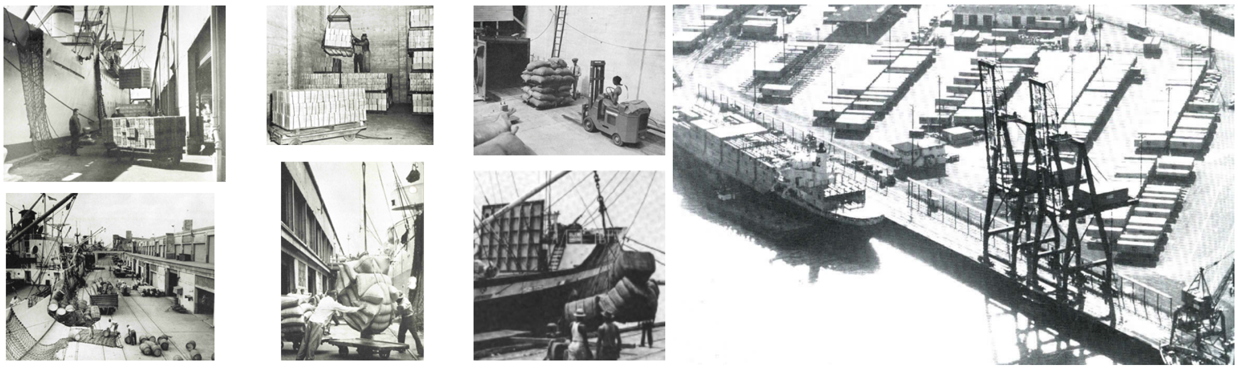

Breakbulk shipping, 1950s Container shipping, 1967

Figure E.1: Illustration of changes in port technology

Notes: Sources: Annual reports for the Port of Seattle and the Port of New Orleans (1950, 1951, 1952, 1955).

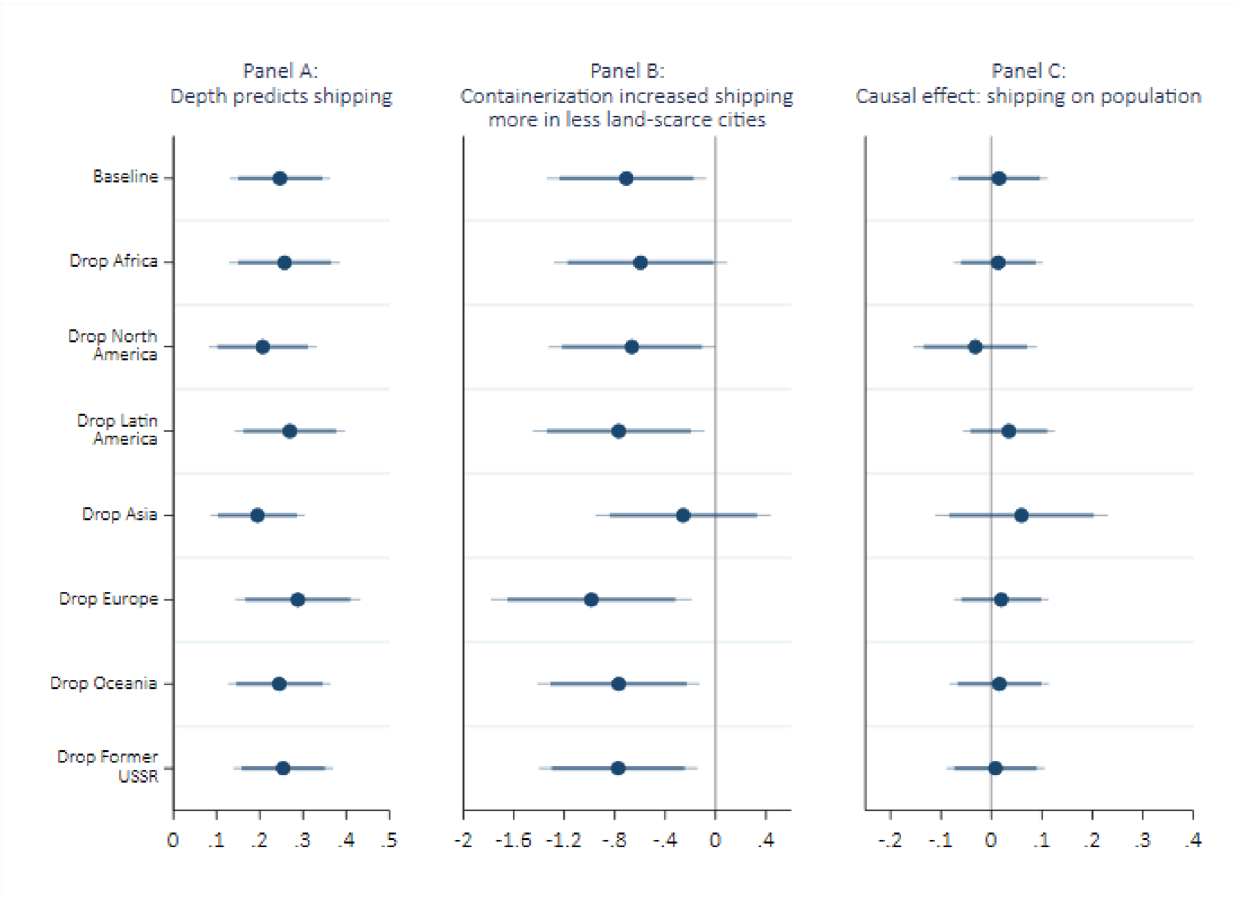

Figure E.2: Dropping continents one at a time

Notes: The plotted coefficients for Panel A are based on the specification in Table 1, column (5). The plotted coefficients for Panel B are based on the specification in Table D.5, column (2). The plotted coefficients for Panel C are based on the specification in Table 2, column (2). 'Baseline' uses the full sample, while the remaining rows drop continents one at a time as labelled.

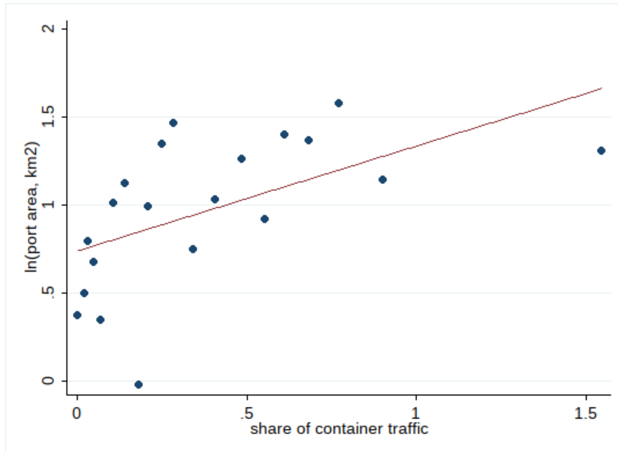

Figure E.3: Relationship between port area and share of container traffic

Notes: The figure shows the correlation between port area and the share of container traffic at the port. The latter is defined as $\frac{\text{container traffic in TEU-s}*12 \text{ tons per TEU}}{\text{total merchandise traffic in tons}}$. As 12 tons per TEU is an approximation of the weight of containerized cargo, it is possible to get a share of container traffic that is larger than one.

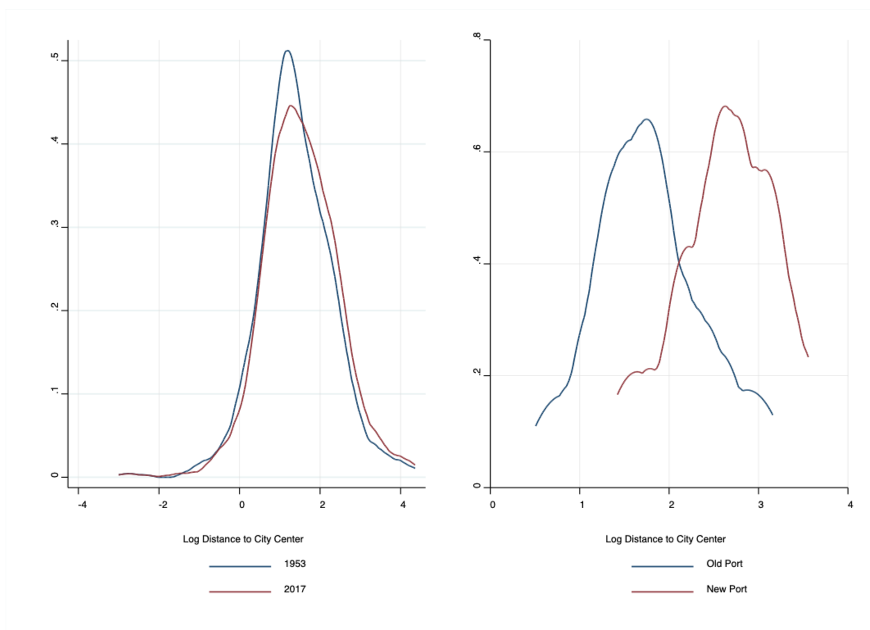

Figure E.4: Location of ports within cities 1953 – 2017

Notes: The figure plots the kernel density of the log distance to the city center for ports in 1953 and 2017. The left panel uses the full sample, the right panel restricts to only those cities where a new port was established after 1953. Data on the geocodes of ports are from the World Port Index. The calculation of city centroids is described in Appendix C.10. Across the full sample, ports moved on average 1 kilometer towards the outskirts (panel A). Where a new port was built, it was on average 9 kilometres further from the centroid of the city than the old port.

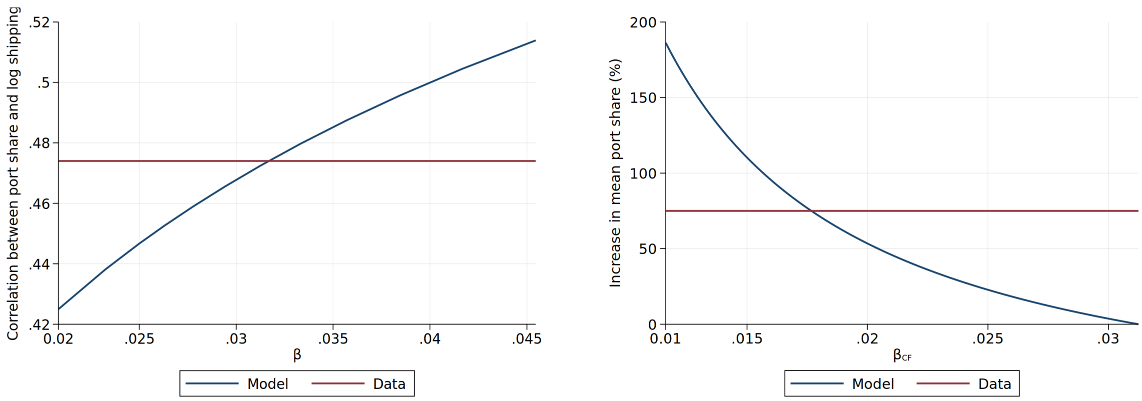

(a) Correlation between port share and shipping as a function of β

(b) Increase in mean port share as a function of β_{CF}

Figure E.5: Calibration of transshipment cost parameter β in 1990 (left) and in the no-containerization counterfactual (right)

Notes: The left panel shows the correlation between the port share and log shipping flows in the model as a function of the transshipment cost parameter β (blue line). It also shows the value of this correlation based on 7 ports in the data (red line). The right panel shows the increase in mean port share in the model, keeping the non-technological determinants of the port share fixed, as a function of β_{CF} (blue line). It also shows the increase in port share induced by containerization in the data (red line).

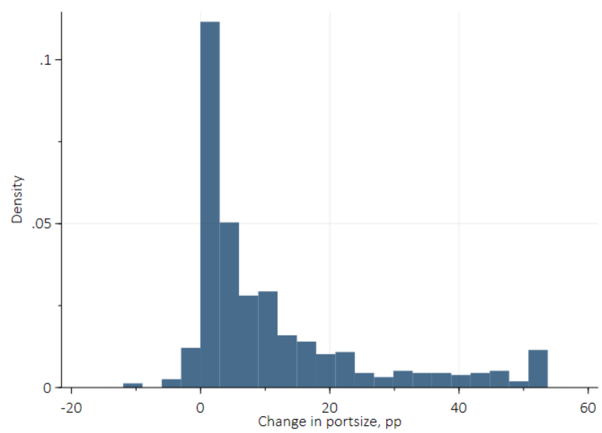

Figure E.6: Histogram of changes in port share between the counterfactual and the 1990 equilibrium, in percentage points

Notes: The figure shows the histogram of the percentage point change in port shares between the model-simulated counterfactual (pre-containerization) and the 1990 equilibrium (after containerization, also model-simulated data).

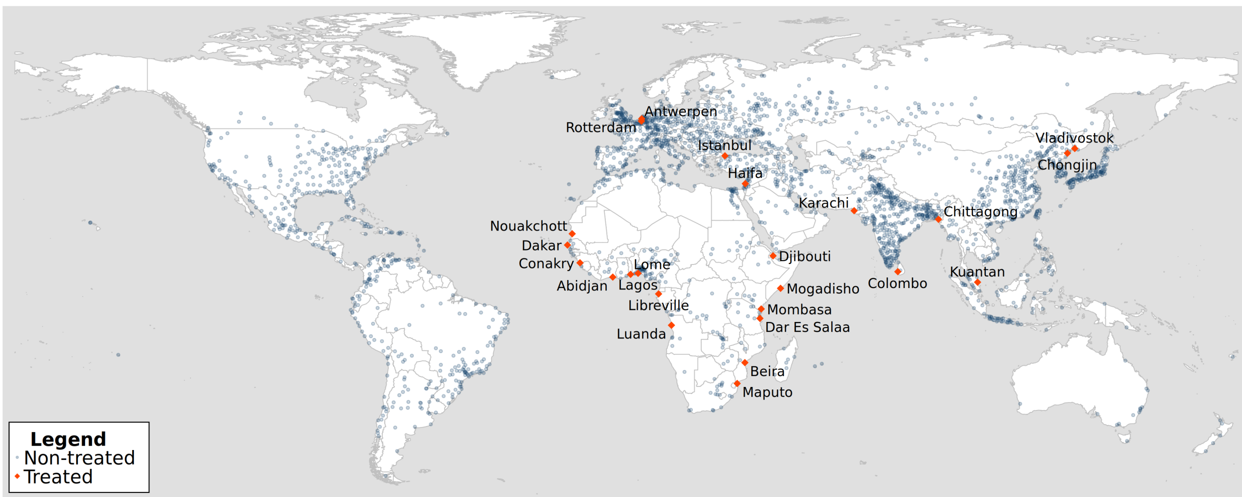

Figure E.7: Maritime Silk Road: targeted ports

Notes: The targeted cities are: Abidjan (Côte d'Ivoire), Antwerpen (Belgium), Beira (Mozambique), Chittagong (Bangladesh), Chongjin (North Korea), Colombo (Sri Lanka), Conakry (Guinea), Dakar (Senegal), Dar Es Salaam (Tanzania), Djibouti (Djibouti), Haifa (Israel), Istanbul (Turkey), Karachi (Pakistan), Kuantan (Malaysia), Lagos (Nigeria), Libreville (Gabon), Lome (Togo), Maputo (Mozambique), Mogadisho (Somalia), Mombasa (Kenya), Nouakchott (Mauritania), Rotterdam (Netherlands) and Vladivostok (Russia). Source: OECD (2018).

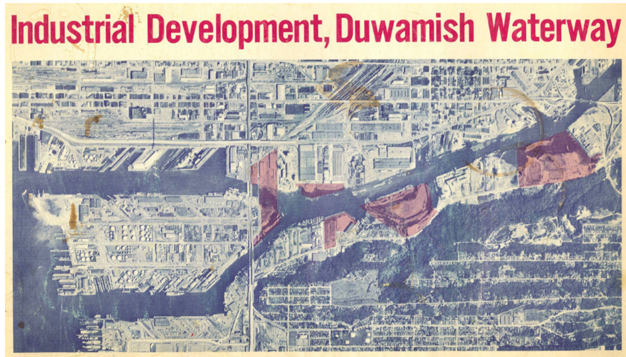

Acquired land parcels (red shading), 1963

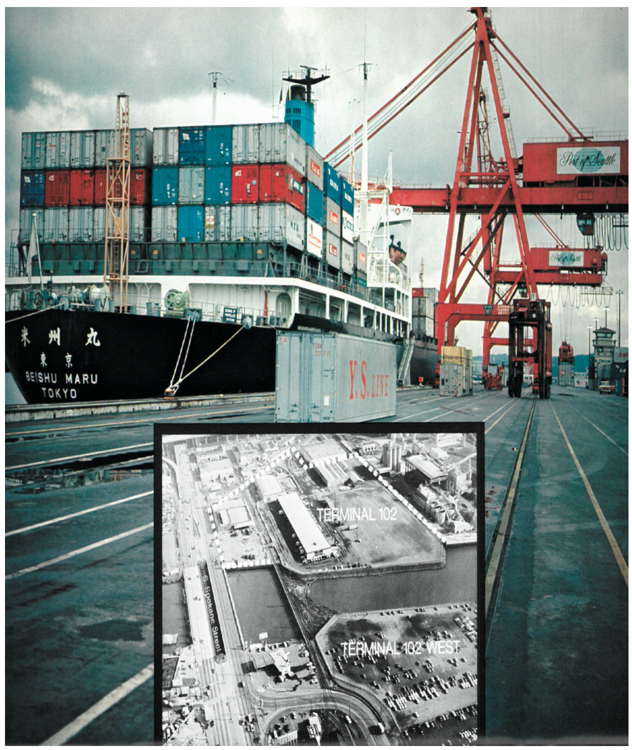

Completed terminal 102, 1970

Figure E.8: Illustration of port development, Seattle

Notes: The two panels illustrate development of the port through the 1960s. The first panel shows the initial set of land parcels acquired by the port along the Duwamish Waterway in the early 1960s. The second shows a container terminal completed in 1970 within this project. Sources: 'Port of Seattle: Industrial Development, Duwamish Waterway' (1963), 'Annual Report of the Port of Seattle' (1970).

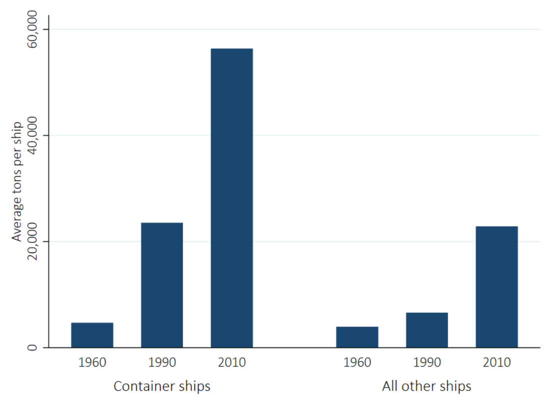

Figure E.9: Development of ship sizes over time, 1960-2010

Notes: The figure illustrates the growth in ship size, as measured in average tons per newly built ship in a given year, for the years 1960, 1990, and 2010, for container-ships and all other ships (i.e., excluding container-ships), respectively. Source: Haworth (2020).

References

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

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

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

Australian Competition and Consumer Commission (2017). Container Stevedoring and Monitoring Report 2016-17. Technical report, ACCC.

Baragwanath, K., R. Goldblatt, G. Hanson, and A. K. Khandelwal (2019). Detecting Urban Markets with Satellite Imagery: An Application to India. *Journal of Urban Economics*, 103–173.

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

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

Borusyak, K. and P. Hull (2022). Non-Random Exposure to Exogenous Shocks. *National Bureau of Economic Research* (Working Paper 27845).

Bridgman, B. (2021). Why Containerization Didn't Reduce Ocean Shipping Costs, At First. *mimeo*.

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

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

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

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

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

Kirchner, M. (2006). Container Vessels and Risk Aggregation: The Cargo Underwriter's View. Technical report, AXA Corporate Solutions.

Lloyd's of London Press (1990). Lloyd's Ports of the World.

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

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

New York Port Authority (1955). 1955 Annual Report. Technical report, New York Port Authority.

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

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

Port Authority of Seattle (1989). Seaport Properties Book. Technical report, Port Authority of Seattle.

Port of Long Beach (1971). Harbor Highlights 1971 Report. Technical report, Port Authority of Long Beach.

Port of New Orleans (1951). 55th Annual Report. Technical report, Port of New Orleans Board of

Commissioners.
Port of New Orleans (1984). 88th Annual Report. Technical report, Port of New Orleans Board of Commissioners.
Port of San Francisco (1971). San Francisco Port Needs, Shipping and Area Requirements. Technical report, Port of San Francisco.
Riffenburgh, R. (2012). A Project History of the Port of Long Beach 1970 to 2010. Technical report, Port of Long Beach.
Saiz, A. (2010). The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics 125*(3), 1253–1296.
UNCTAD (1971). Review of Maritime Transport. Technical report, United Nations Publication.