Working Paper

The Local Impact of Containerization

Leah Brooks
Trachtenberg School of Public Policy and Public Administration
George Washington University

Nicolas Gendron-Carrier
Department of Economics
University of Toronto

Gisela Rua
Division of Research and Statistics
Federal Reserve Board of Governors

Monday 2nd November, 2015
Preliminary. Please do not cite without authors' permission.

We are very appreciative to Matthew Turner and Gilles Duranton for sharing the 1956 County Business Patterns data. We are also very appreciative to research assistants Daniel Walker, Adrian Hamins-Puertolas, Arthi Rabbane and Alex Severn who helped with data entry. We thank Elliot Anenberg, Paul Carrillo, Marco Gonzalez-Navarro, Justin Pierce, Daniel Trefler, and Matthew Turner for many helpful comments and discussions. Finally, we thank Nora Brooks, retired reference librarian and Leah's mother, whose combination of digital and pre-digital skills uncovered the 1948 port-level trade data.

The Local Impact of Containerization

future abstract here!

Leah Brooks
Trachtenberg School of Public Policy
George Washington University
805 21st St. NW, Room 601F
Washington, DC 20006
lfbrooks@gwu.edu

Nicolas Gendron-Carrier
Department of Economics
University of Toronto
150 St. George Street
Toronto, ON, Canada, M5S 3G7
nicolas.gendron.carrier@mail.utoronto.ca

Gisela Rua
Division of Research and Statistics
Federal Reserve Board of Governors
20th and C Streets, Mail Stop 82
Washington, DC 20551
gisela.rua@frb.gov

1 Introduction

Despite a vast and prominent theoretical literature that emphasizes the role of international market access in explaining variation in economic activity (e.g. Krugman, 1991; Fujita et al., 1999), we know very little empirically about the causal effects of trade on local economic growth. Trade's impact on the distribution of population is of central interest, since dense population clusters/cities are the key location for innovation and the creation of human capital (Lucas, 1988). In this paper, we use the advent of containerization—a technological shock that dramatically reduced international shipping costs—to examine how access to international markets affects city growth.

Containerization is premised on a simple insight: Packaging goods for waterborne trade into a standardized container makes them cheaper to move. Containerization simplifies packing, transit, pricing, and the transfer from ship to train to truck; it also limits previously frequent and lucrative pilferage. Since the advent of containerization in 1956, international trade has grown tremendously.[1] Bernhofen et al. (forthcoming) estimate that containerization caused international trade to grow by more than 1000 percent in the 15 years since 1966. Containerized trade now dominates ocean shipping, and containers account for well over 75 percent of U.S. domestic rail traffic (Rodrigue, 2015).

Containerization's impact on cities is theoretically ambiguous. The new economic geography literature predicts that firms will locate in regions where market potential is high because these regions are the most profitable. A decline in international trade costs can promote local economic growth by allowing firms in the region to reach foreign consumers, thereby attracting new firms to the region and increasing local employment

---

[1] There is some dispute about the magnitude of the cost decline caused by containerization. While Bridgman (2014) and Hummels (2007) argue that direct declines in shipping rates were small, Levinson (2008) argues that declines were large. We address this issue in Section 2.

(Ramondo et al., 2014; Redding, 2015). Improved access to international markets may, however, make a region poorer if increased foreign competition causes firms to depart the region entirely for a lower input area, for example China (Autor et al., 2013). In addition, it may increase land and congestion costs which cause firms to relocate. Our empirical work tests whether, in response to this sharp decline in international transportation costs, the agglomerative forces keeping firms and workers in cities are stronger than those repelling them.

Our unit of analysis is the city. We therefore require data describing the evolution of population and port facilities for a panel of cities. We construct two such panels for 1950–2010: one for the world and one for the United States. We combine these data with port level data on location, depth, and size. For the U.S., we additionally include measures of pre-containerization international trade by port.

We use these data in both reduced form and general equilibrium approaches. In the reduced form work, we address the non-random selection of cities into proximity to a containerized port with a novel instrument—specifically, city proximity to a very deep port in 1953 is an instrument for city proximity to a containerized port.

The first requirement for a good instrument is that it is correlated with the endogenous variable. Container ships are substantially larger than their predecessors, and displace more water. They therefore require deep ports. While a port can be arbitrarily deep in the absence of cost concerns and environmental regulations, initially deeper ports are cheaper to convert to container ports because they require less drilling and dredging. This instrument is analogous to the supply shifter instruments used in the industrial organization literature. Empirically, we find a very strong relationship between the instrument and the endogenous variable.

A good instrument must also impact city population only through its relationship with a city's proximity to a containerized port. Although ports varied in depth before containerization, being a very deep port—beyond 25 or 30 feet—posed no particular competitive advantage. Most ships did not displace enough water to require more depth. Crucially, being a very deep port matters only after the invention and diffusion of containerized shipping. Thus, we paramaterize our instrument as a city's proximity to a very deep port, where the depth cut-off is beyond what was generally considered a useful depth in the pre-containerization era.

Our causal estimates of the impact of containerization on the growth of cities rely on the quasi-random variation in initial depths. The estimates compare cities that are treated with a container terminal—because they had nearby ports that were very deep before the invention of containerization—to otherwise similar cities.

In our world panel, we find that, from 1950 to 2010, cities within 100 km of a containerized port grew about 25 percentage points more than other cities. This is 16 percent of the mean city population growth over the period. Effects are strongest for cities within 100 kilometers of a containerized port, and are marginally statistically significant at distances greater than 200 kilometers. As with our world sample, results for the U.S. show that proximity to a containerized port causes significant population growth. Our most complete instrumental variables estimates indicate that being a county 0 to 50 km from a container port is associated with a 30 percent increase in population growth. In the U.S., we also find statistically significant population growth for cities somewhat farther from containerized ports: Population growth in counties 200 to 250 km from containerized ports experience changes equal to 35 percent of the mean growth over the period.

These reduced form methods help us understand the shift in the distribution of population. However, these methods do not allow us to assess whether population increases near ports come at the expense of other locations. To tackle this general equilibrium proposition, we turn to a market access analysis, as in Donaldson and Hornbeck (forthcoming). We are currently assembling historical maps for this purpose, and anticipate having the results in a future draft.

Our paper is closely related to previous work that examines the role of market access in explaining spatial variation in economic activity (e.g. Davis and Weinstein, 2002; Hanson, 2005; Redding and Sturm, 2008). These authors consider a variety of changes in market access, ranging from the bombing of Japanese cities during the Second World War to the division and reunification of Germany, to test new economic geography predictions. Our paper contributes to this literature by considering directly the effects of a large decline in international transportation costs on the growth of cities.[2]

There is also an active academic literature investigating the effect of transportation infrastructure on the growth of cities (e.g. Michaels, 2008; Duranton and Turner, 2012; Donaldson, forthcoming; Gonzalez-Navarro and Turner, 2015). These studies examine how investments in highways, railways, and subways, have shaped regions and cities.[3] Most existing studies in this vein consider infrastructure that reduce domestic transportation costs. Our results contribute to this literature by showing that investments in transportation infrastructure that reduce international transportation costs, such as the construction of new container terminals, can also improve the economic condition of target areas.

Finally, our work draws on the large literature concerned with the fundamental determinants of economic growth, pioneered by the work of Barro (1991). A consistent finding in this literature is that landlocked countries are much poorer than other countries. Our results lend credence to the hypothesis that good access to international markets matters for economic growth.[4]

---

[2] Our paper is also related to a growing literature in international trade that looks at the impact of trade on local labour markets (e.g. Topalova, 2010; Autor et al., 2013; Kondo, 2013; Dix-Carneiro and Kovak, 2015). These studies suggest that trade can have substantial localized effects.

[3] See Redding and Turner (2015) for a survey of the literature.

[4] Romer and Frankel (1999) Feyrer (2009), and Pascali (2014) provide country level evidence of the effects of trade on economic growth.

The remainder of this paper is organized as follows. The following section provides background on containerization, and Section 3 discusses the data. We present empirical methods in Section 4, and results in Section 5. We conclude with Section 6.

2 Containerization

Before the advent of containerization, shipping was expensive and slow. Vessels spent weeks at ports while cargo was handled, piece by piece, by gangs of dockworkers. Containerization brought about an unprecedented change in transportation: It made moving goods across the world dramatically easier, cheaper, and faster. Since Malcolm McLean's first application of containerization to sea and land transportation in 1956, there was a tremendous growth in international trade. Bernhofen et al. (forthcoming) estimate that containerization caused international trade to grow by more than 1000 percent over the 15 years following 1966.

Containerized shipping began in the United States in the mid-1950s. Figure 1b reports the total number of container ports per year in the United States. As we can see, the bulk of containerization adoption occurred in the 1960s and 1970s, while a smaller number of adoptions occurred in the 1980s, and an even less after that. Adoption of containerization was exceptionally rapid, not only within the United States, but also across the world (Rua, 2014). Figure 1a shows the total number of container ports worldwide by year, with the bulk of international adoption occuring in the 1970s. As of 2013, containerized trade accounted for over half of global commodity trade (United Nations Conference on Trade and Development, 2013).[5]

Containerization's success relies on two key innovations. The first is the mechaniza-

---

[5] While containers are appropriate for carrying many goods, as diverse as toys and frozen meat, some goods are not yet containerizable. Both "non-dry cargo" and "dry-bulk commodities" such as oil, fertilizers, ore, and grain cannot be shipped inside "the box."

tion of container movements in and out of ships and around the port, using specialized container cranes. By lifting containers onto rail cars and trucks, container cranes also ease transportation beyond the port. This simple technological innovation radically changed the entire process of on- and off-loading: Instead of spending weeks at ports, like in the breakbulk/pre-containerization era, container ships can now spend just a few days. And this quicker turnaround makes larger ships more profitable. Moreover, since containerization greatly reduces the risk of loss and damage, it allows all different kinds of goods, with different destinations, to be shipped together.[6]

The second key innovation is the development of common standards for container size, stacking techniques, and grip mechanisms. These international standards allow a container to be used across modes of transportation—ships, trucks, rail—and across countries. While the U.S. standard for containers was adopted in the early 1960s, the international standard, of the International Organization for Standardization, was promulgated in the late 1960s. These national and international agreements on standards can be viewed as a successful resolution for a potentially severe collective action problem. (In Figure 1, we note the dates of the U.S. and international standards adoption, as well as when international diffusion plateaus in the early 1980s (Rua, 2014).)

A strand of the trade literature argues that containerization did little to lower direct ocean shipping costs. However, as Hummels (2007) explains, direct shipping costs do not full capture containerization's impact on the quality of shipping services, in particular delivery time and pilferage and damage risks. To the extent that reductions in ocean shipping times and in pilferage and damage risks do not show up in traditional measures of shipping costs, those measures fail to capture containerization's full impact on transportation costs. In fact, Hummels and Schaur (2013) estimate that each transit

---

[6] Losses to pilferage have plummeted. Wilson (1982) estimates loses to pilferage at roughly 25% in the breakbulk era, and near zero in the container era.

day is worth between 0.6 to 2.1 percent of the value of the good, lending credence to the argument that unmeasured benefits of containerized shipping are non-negligible.

For the purposes of this paper, and consistent with the industry definition, we call a port "containerized" when it has special infrastructure and equipment to handle containers. Specifically, the port has invested in equipment to handle shipping containers which enables their movement in and out of ship and onto a train or a truck. Container ports also require extensive marshalling yard in which containers in transit can wait to be moved (Rua, 2014).

The conversion from a traditional port to a containerized port is exceedingly expensive. For example, Kendall (1986) writes that "In the period between 1968 and 1973, shipowners, terminal operators, and port agencies in the United States alone invested seven and a half billion dollars in ships, containers, and port facilities." In addition, a survey published by the *The Journal of Commerce and Commercial* in 1977 estimated that acquiring container-handling cranes cost \$1.75 million per 30-ton capacity and securing waterfront space for terminals and marshaling yards cost \$250,000-\$300,000 per acre (Morison, 1977).[7]

3 Data

To study containerization's impact on cities, we require data describing the evolution of population and port facilities for a panel of cities. We construct two such panels: one for the world and one for the U.S. This section summarizes the data; full details are in the data appendix.

Throughout, our unit of observation is a city in a year. More precisely, cities are urban agglomerations in the international data, and counties for the U.S. sample. We

---

[7] See Talley (2002) for containerization's impact on dockworkers.

frequently use the term city for both for expositional ease.

Our sample frame for international cities is the 2014 Revision of World Urbanization Prospects. This dataset contains all 1,692 urban agglomerations with population exceeding 300,000 at any time between 1950 and 2014. By construction, this sample over-represents fast growing cities that were small in 1950 but grew rapidly in the second half of the twentieth century. To mitigate this sampling issue, we also restrict the sample to cities with population over 50,000 in 1950, which yields a world panel of 1,051 cities. Our results are robust to different sample selection criteria.

For the United States, our sample frame is the county level Decennial Census, for years 1910–2010.[8] We assemble a time-invariant panel of counties by aggregating 1950 counties to their 2010 counterparts (most county changes 1950 to 2010 are splits) and by dropping a very few counties with land area changes greater than 35 percent. For the period of analysis, 1950 to 2010, we observe population, employment, share of manufacturing employment, age distribution, income distribution, and education by county in each year.[9] We omit Alaska from our analysis because its administrative districts in 1950 do not correspond to its modern counties. This yields 2,702 counties with complete data, compared to the 2010 total of 3,007 counties.

To these sample frames, we add port attribute data. Our universe of ports is all ports that existed in either 1953 or 2015, as defined by the 1953 and 2015 *World Port Index*.[10] For each port, we observe its location (latitude and longitude), size (in 4 discrete categories), and depth (in 8 discrete categories). We use depth of the wharf in 1953 as our measure

---

[8] For the 2010 sample, we use the Decennial Census for population figures and the American Community Survey (years 2008–2012) for other demographic covariates.

[9] The share of manufacturing employment comes either from the Decennial Census or from the 1956 County Business Patterns, which we received courtesy of Matt Turner and Gilles Duranton. We hope to expand our use of these data in future drafts, using the 1956 version as our pre-containerization period.

[10] At this point, for the world analysis we only use 1953 ports that were classified as larger than "very small" in 2015. We are currently entering data to be able to use the sample of all 1953 ports, regardless of size.

of pre-containerization port depth.[11] The year of first containerization comes from the *Containerisation International Yearbook*, volumes 1968 and 1970 to 2010. For the U.S. only, we observe 1948 and 1955 international trade in dollars by port from the Census Bureau's Foreign Trade Statistics.

We associate each city with a vector of ports and port characteristics. First, we calculate the distance from each city center to each port. For the world sample, the World Urbanization Prospects data gives us the latitudes and longitudes of city centers; for the U.S. sample, we use the geographic center of the (grouped) county. Second, for each distance bin from city center $d \in D$, we calculate the number of 1953 ports, the year of first containerization across all ports, and the maximum 1953 depth across all ports.[12] For the U.S., we also calculate total international trade volume in 1948 and 1955 for all ports in each distance bin.

The worldwide sample has the benefit of describing a larger share of cities affected by changes in international trade, but offers a truncated view of containerization's impact on population, since cities must be of sufficient size to enter the sample. In contrast, the U.S. sample describes a much smaller slice of containerized trade, and one that is clearly not typical, given that the U.S. was the first country to have large-scale adoption of containerization. However, because the U.S. sample covers all continental U.S. land area, it does allow for a complete description of containerization's impact on the spatial organization of economic activity. In addition, the U.S. sample allows for the inclusion

---

[11] Results are generally robust to using anchorage and channel depth, which the *World Port Index* also reports.

[12] Appendix Figure 1 is a graphical depiction of this classification for counties in Southern California. Red triangles mark the geographic center of the county, called a centroid. Containerized ports are denoted with the pink anchor, and the grey circles show rings of 50 kilometers around each port (the Los Angeles and Long Beach ports are next door to one another). Only one county centroid is within the 50 kilometer ring. Thus, at the distance interval of 0 to 50 kilometers, only Orange County is treated with having a containerized port in this range. The bottom panel shows that at a distance of between 50 and 100 km from a port, both San Diego (the southernmost) and Los Angeles are treated counties. Ventura County, just to the north of Los Angeles, remains untreated.

of more detailed covariates.

Table 1 reports summary statistics for both samples, using the format we follow throughout the paper: Panel (a) presents world statistics, and panel (b) presents U.S. statistics. We report summary statistics for six distance-to-containerized-port bins (columns 1 to 6), and by the categories of ever and never containerized (columns 7 and 8). A city may appear in more than one distance-to-containerized-port bin, but the number of observations in the "ever" and "never" columns sums up to the total sample size.

The worldwide panel of Table 1 reports log population—our main dependent variable—in 1950, the pre-containerization year, and 2010, the final year of the sample. In 1950, cities near future containerized ports are roughly 30 percent larger than cities never near a containerized port (columns 7 and 8). These differences between cities with and without container ports generate a possible bias in the OLS estimation that we address in the empirical section. In addition, cities closer to future containerized ports are larger than those farther from containerized ports: 1950 log population declines from column 1 to column 6, with only small exceptions. Comparing 1950 log population to 2010 log population, cities farther from container ports (column 6) have, on average, smaller population increases than cities closer to container ports (column 1).

The second section of the table shows that about fifty percent of cities in the sample are in Asia, roughly one-fifth are in Europe, a slightly smaller fraction are in North America, and the remainder of cities are split between South America, Australia, and Africa. Also, the average container port has existed for 35 years (column 7).

Table 1b repeats the analysis for the U.S. sample. We observe U.S. population every ten years from 1910 to 2010. As in the worldwide sample, from 1910 to 1950, log population in cities near future container ports is larger, and it increases at a faster rate, than in cities farther from future container ports. Not surprisingly, given that the U.S. sample covers the entire country, the average population among ever-containerized U.S. cities is almost 2 log points (or 200,000 people) smaller than among ever-containerized world cities.

Using the Census Bureau's division of the U.S. into four regions, we find that almost half of the counties within 300 km of container ports are in the Southern region, slightly under one-third are in the Midwest, about one-sixth are in the Northeast, and nearly 1 percent are in the West (note that Western counties tend to be geographically larger). On average, cities near future container ports tend to have substantially more of their employed population in the manufacturing sector than cities never near future container ports—43 versus 27 percent (columns 7 and 8).

Finally, unsurprisingly, the average U.S. city has had longer exposure to containerization than the average world city. U.S. cities near container ports have been, on average, near container ports for 44 years (column 7), relative to 35 years for the average world city. Appendix Table 1 reports total 1948 and 1955 international trade by distance bin to city and shows that cities near future container ports have, on average, more pre-containerization international waterborne trade.

4 Empirical Methods

In this section, we explain our empirical strategy for estimating the causal effect of being near a container port on city population growth. We look at city population because it can be interpreted as a summary statistic for overall welfare.

We begin by presenting a naive regression of containerization's impact on population to clearly explain the potential endogeneity issues. We then motivate and explain our instrumental variable strategy. We conclude by discussing how we paramaterize a city's proximity to a containerized port.

4.1 First Difference Specification

Our goal is to understand how city population responds to the advent of containerization. Empirically, we measure this in two ways: whether a city is near a port that ever containerizes, and the number of years the city has been near a containerized port. We focus our discussion in this section on the former measure, which is easier to interpret, but we discuss both measures in the estimation results section.

We estimate an equation of the form:

$$\Delta y_{i,t} = \beta_0 + \beta_1 \Delta c_{i,t} + \beta_2 x_{i,1950} + \Delta \epsilon_{i,t} \ , \tag{1}$$

where $i \in I$ are cities, and $t \in T$ are years. Our dependent variable, population, is $y_{i,t}$. The operator Δ denotes long-run differences, so that $\Delta y_{i,t} = y_{i,t} - y_{i,1950}$. Since there were no container ports in 1950 ($c_{i,1950} = 0 \ \forall i$), our main explanatory variable of interest is whether a city is near a container port at time t, $c_{i,t}$. We also control for baseline covariates in $x_{i,1950}$. Standard errors are clustered at the city level, which is equivalent to having robust standard errors in this two-period case.

To establish the causal effects of containerization on the growth of cities, however, we must contend with the selection of cities that are proximate to container ports. For example, if economically healthier cities, which are more likely to take advantage of increased trade, are also more likely to be proximate to container ports, OLS estimates are biased upward. In contrast, if large cities attract proximate container ports, to take advantage of the larger markets those cities offer, our estimates of containerization's impact on city growth would be biased downward. This is because large cities, on average, grow more slowly than smaller cities.

The first difference strategy addresses some of these concerns. The first difference nets out time-invariant city characteristics that may make container terminals more likely to locate near particular cities. For example, equation 1 controls for changes in population due to a city's geographic location, its long-run industry mix, or its climate. This method also accounts for changes in population that impact all cities equally from 1950 to 2010, for example an economic downturn—worldwide or in the U.S.—that might have impacted the likelihood of containerization adoption.

In addition, in a first difference approach, in contrast to a panel fixed effects approach, we can also control for initial conditions, $x_{i,1950}$. Including initial conditions in the first difference model allows for powerful controls, such as differential trends in population growth by initial period covariates. Therefore, we can directly address the concern of differential growth rates by initial city size. For both the world and U.S. samples, we further control for being within 300 km of a 1953 port, the number of 1953 ports within 300 km, and the initial population in 1950. For the U.S. sample, we also control for population in 1920 to 1940, the 1956 manufacturing share of employment, and the total value of 1955 international trade.

Nevertheless, these estimates do not allow a distinction between population reallocation or net growth (Redding and Turner, 2015). A positive estimate for β_1 could result from a mix of domestic migration, international migration, or natural population increase. In future work, we hope to be able to dissect some of these differences in the U.S. data.

This empirical strategy would yield a causal estimate for the effect of containerization on population if containerization were exogenous, conditional on time-invariant factors at the city level and on the initial covariates that we include. However, suppose that cities are more likely to be near a container port if they made better industrial choices in the 1940s and early 1950s. This is something we may fail to capture, even after netting out city-specific time-invariant factors. This type of endogeneity would yield a positive bias in the OLS estimates. Conversely, if there is more containerization adoption near cities with less successful industrial choices in the 1940s and 1950s, the OLS estimates would be biased downwards.

4.2 Instrumental Variables

To deal with selection bias in the adoption of containerization, we use being near a very deep port in 1953, z_i, as an instrument for containerization $\Delta c_{i,t}$:

$$\Delta c_{i,t} = \alpha_0 + \alpha_1 z_i + \alpha_2 x_{i,1950} + \eta_i \ , \tag{2}$$

There are two requirements for a successful instrumental variable strategy. The first is a strong relationship between containerization and initial depth. The second requirement is that, conditional on covariates, being near a very deep 1953 port is uncorrelated with unobserved determinants of population growth between 1950 and period t. In other words, proximity to a very deep 1953 port affects city i's population growth only through its impact on containerization:

$$\text{Cov}(z_i, \Delta\epsilon_{i,t}) = 0 \tag{3}$$

Conditional on these assumptions, β_1 yields a causal estimate of proximity to a container port on population growth.

We explore the two instrumental variable requirements in turn. First, we anticipate that city proximity to a very deep port pre-containerization should be strongly correlated to city proximity to a container port. Even the first container ships were substantially larger than their predecessors, and larger ships sit deeper in the water and require greater depth to navigate and dock.

Although harbor depth is malleable, it is malleable only at great cost. Given enough money and sufficiently lax environmental regulation, a harbor can arguably be made arbitrarily deep. However, ports which are initially deep have a competitive advantage when technology changes to favor very deep ports. This inability of all ports to adjust equally is confirmed by Broeze, who notes that while "ship designers [keep] turning out larger and larger vessels," and "the engineering limits of port construction and channel deepening have by no means been reached[, t]his, however, may not be said of the capacity of all port authorities to carry the cost of such ventures" Broeze (2002, pp. 175–177). Converting a breakbulk port into a container port is substantially cheaper when the harbor is already deep.

This intuition is borne out in practice by containerization adoption patterns. Figure 2 shows how the likelihood of a city being within 300 km of a containerized port varies over time with proximity (within 300 km) to ports of a given depth. Panel (a) is based on worldwide cities, and panel (b) is based on U.S. counties. In both cases, we see a strong relationship between proximity to deep 1953 ports and later proximity to containerized ports. In the world sample, cities are more likely to be near a container port earlier when they are near a deep 1953 port. By 1975, virtually all of cities within 300 km of a port that was greater than 40-feet deep in 1953 are near a container port (blue line). Adoption is substantially slower for international cities near ports that are less than 20-feet deep, though roughly eighty percent of these cities are near a containerized port by the end of the sample period.

The U.S. results in panel (b) are very similar, but they show an even more-marked pattern of county proximity to a containerization port by depth.[13] Counties within 300 km of a port deeper than 35 feet are virtually always within 300 km of a container port by the end of the sample period. Only roughly 25 percent of counties within 300 km of

---

[13] This may in part be due to the fact that we use all ports, not just all not-very-small 2014 ports, in the U.S. analysis. We plan to update the world sample to include all 1953 ports in our next draft.

ports with depths between 25 and 35 feet are not near a container port by the end of the sample period. For counties near less deep ports, however, containerization is decidedly not a certainty. Indeed, counties that are near only shallow ports—those less than 20 feet deep—are never near a container port.

In Table 2, we show that these differences are statistically meaningful. The specification in the table is not precisely what we will use in our estimation, which relies on a vector of port proximity measures and a vector of depth proximity measures as instruments, but it illustrates clearly the intuition behind our method. We defer discussion of the precise instrument specifications to the end of this section, and results in the following section.

The first column of Table 2a shows that, relative to cities near ports of 1953 maximal depth less than 10 feet, cities near ports of 1953 maximal depth 40 feet or above are more than fifty percent more likely to be near a container port in 2010. In the world panel, this likelihood is never less than fifty percent for cities near ports greater than 30 feet deep in 1953, and it is always significant at the five percent level. For cities near ports 30 feet deep or less, estimated coefficients are under fifty percent and less likely to be significant. This motivates our specification in the second column, where we use an indicator variable for proximity to ports of a maximal 1953 depth of 30 feet or more. Here, proximity to a very deep port makes a city 10 percent more likely to be within 300 kilometers of a container port. The F statistic for this estimation is 10, indicating a reasonably strong relationship between the instrument and the endogenous variable.

Columns 3 and 4 repeat the same specifications, but the dependent variable is a city's years since the first containerization across all ports within 300 kilometers. Consistent with Figure 2a, we see the strong pattern of cities near deeper ports having had longer exposure to containerization. Relative to cities near ports less-than-10 feet deep, cities near ports 40 or more feet deep have had containerization for 18 more years. According to Table 1a, the average city within 300 kilometers of a container port has had 35 years of containerization exposure, so the depth measure explains a large portion of the variation. This explanatory power falls, as we expect it should, with depth, and is uniformly insignificant for depths below 30 feet. Using the dichotomous specification we will rely on for our estimates, column 4 shows that cities near ports that are 30 or more feet deep experience an additional 9 and a half years of exposure to containerization. The F statistic for this specification is a very robust 55.

Table 2b repeats the same specification for the U.S. sample. Here, the relationship between proximity to deep 1953 ports and later proximity to container ports is even more striking. Cities near ports that are over 40 feet deep in 1953 are certain to be near containerized ports in 2010 (relative to cities near ports of 1953 depth less than 10 feet), and the coefficient declines almost monotonically with depth. Consistent with Figure 2b, cities near ports that are less than 20 feet deep in 1953 are very unlikely to be near a container port in 2010; both coefficients are near zero and insignificant. Using the dichotomous specification, cities near ports greater than 30 feet deep in 1953 are 7 percent more likely to be near a container port in 2010. This specification has a robust F statistic of 41.

As before, in the final two columns of the table, the dependent variable is a city's years of exposure to containerization. Cities near ports that are 40 feet or more deep experience about 20 years more exposure since first containerization. This coefficient falls monotonically with depth, consistent with our argument about how depth should impact the likelihood of containerization. Using the dichotomous specification, proximity to a 1953 port that is 30 or more feet deep increases a city's exposure to containerization by almost four additional years, or about ten percent of the mean from Table 1b.

Given this evidence, which is consistent with a strong relationship between the dependent variable and the instrument, we now turn to the second condition for instrument validity—that proximity to a very deep 1953 port affects city i's population growth only through its impact on containerization. A key concern with the instrument is that port depth may explain city success even before containerization. This is surely true. Being a port deep enough for pre-containerization ships undoubtedly helped to generate cities near ports. However, the minimum depth for pre-containerization success was substantially shallower than the minimum depth most useful for containerization. We account for this concern by limiting the depth variation in the instrument to be binary: whether the port is very deep in 1953.

Before containerization, port depth conveyed some advantage, but it was not particularly useful for a port to be very deep. Given the limited draft of breakbulk ships, greater depth was only useful up to a certain point. This is clear even from how data on port depth was collected. The 1953 *World Port Index*'s deepest category is "40 feet and above," while the deepest category in the 2015 *World Port Index* is "76 feet and over."

Our claim that depths beyond 30 feet were not particularly advantageous to port success is supported by a number of contemporary commentators. As late as 1952, F. W. Morgan argues in *Ports and Harbours* that beyond a certain level, depth is not a particularly useful feature of a port:

> The importance for a few ports of maintaining a ruling depth sufficient to admit the largest liners [a draft of 40 feet] emphasizes unduly their importance to the port world. A super-liner which comes into a port every few weeks will, it is true, amplify that port's tonnage figures by half a million tons or so annually. . . . The greater part of world trade by sea and the greater part of the traffic of many ports is concerned with ships of more modest size.
>
> It would certainly be possible to devise a classification of ports by the draught of ship which can be berthed in them. Halifax and Wellington would appear in the first class, and their ability to berth the largest ships is a great asset

> in wartime. It tells, however, only a little about their normal significance as ports. (p. 15, Morgan (1952))

Earlier writers also confirm this view. A 1938 monograph argues that "For the ports with which we are dealing, the 30-foot channel at low-water will be taken as the minimum standard in relation to the needs of modern ships" (Sargent, 1938).[14] However, he notes that the cost of making a channel deeper is no small endeavor: "It is a question how far the rest of the world, Europe in particular, is prepared, except in special circumstances, to face the very heavy cost of providing for the needs of the ocean mammoth" (Sargent, 1938, p. 21).

Thus, our instrument is analogous to a cost shifting instrument in the industrial organization literature. If our instrument functions as a price shifter—shifting the supply of ports after the advent of containerization, but not the demand for ports—it should be unrelated to port demand.

In Section 5, we empirically allay concerns that the instrument is correlated with pre-containerization changes at the city level. To do so, we examine the correlation between pre-containerization factors and the identifying variation in the instrument.

4.3 Paramaterization of Distance to Container Port

Until now, we have treated proximity to a containerized port as a uniform category. In practice, our specification allows for different impacts of proximity to a container port on population growth, by distance to a container port. Therefore, we measure change in access to container ports by:

$$\Delta c_{i,t} \equiv \sum_{d \in D} \beta_{1,d} \mathbb{1}\{\text{Container port between } d_1 \text{ and } d_2 \text{ km}\}_{i,t} \,, \tag{4}$$

---

[14] He goes on to write that in the U.S., a 35-foot draught is becoming standard (p. 21).

where $d \in D$ is a set of distance bins for city i, and $\{d_1, d_2\}$ are the lower and upper bounds of each bin. In the world sample, our bins in kilometers are $\{0-100, 100-200, 200-300\}$. For the U.S. sample, we use kilometer bins of $\{0-50, 50-100, 100-150, 150-200, 200-250, 250-300\}$.

This flexible paramaterization allows for potentially non-linear effects of distance to the container port on population growth. Our goal with this paramaterization is to let the data tell us whether cities need to be very near container ports to experience gains from trade, or whether the specifics of this technology allow for more dispersed growth.

We also propose a set of instrumental variables that parallels this specification:

$$\Delta c_{i,t} = \alpha_0 + \sum_{d \in D} \alpha_{1,d} \mathbb{1}\{\text{Very deep port in 1953 between } d_1 \text{ and } d_2 \text{ km}\}_i + \alpha_2 x_{i,1950} + \eta_i \,, \tag{5}$$

where $\mathbb{1}\{\text{Very deep port in 1953 between } d_1 \text{ and } d_2 \text{ km}\}_i$ is a dummy variable equal to 1 if the maximum depth of any 1953 port in the bin d_1 to d_2 is greater than 30 feet.

5 Results

We now turn to estimates of the impact of proximity to a containerized port on city population growth. We first discuss the world results, starting with OLS, followed by the first-stage instrument results, and then the full two-stage least squares results. We then repeat this pattern for the U.S., where we add additional tests of instrument validity.

5.1 World Results

Table 4a presents OLS results for the relationship between proximity to a container port and population growth from 1950 to 2010. For the world sample, we measure city proximity to container ports in three bins: 0 to 100 km, 100 to 200 km, and 200 to 300 km. These bins are not mutually exclusive: A city can be proximate to a container port in more than one bin.

Column 1 controls only for whether the city is near (within 300 km) of a port in 1953, which allows for differential population growth trends for cities initially near and far from ports. In this specification, cities that are 100 to 200 km and 200 to 300 km from containerized ports see statistically significant declines in population. These coefficient estimates translate to a roughly 5 percent decline in the population growth rate (the dependent variable mean is at the bottom of the table). Column 2 additionally controls for the number of ports within 300 km of the city in 1953, to measure port intensity. With this control, the estimates remain negative, but decline in magnitude and become insignificant.

Column 3 adds country fixed effects, which allow for different trends in population growth by country. Most of the estimated coefficients change signs. We interpret this sign switch as evidence that cities more likely to be near container ports are, on average, in countries with slower population growth than cities less likely to be near container ports. This is consistent with greater adoption of container technology in more developed countries, which have slower rates of overall population growth and slower rates of urban population growth.[15] However, within any given country, cities less than 100 km and within 200 to 300 km of container ports see relatively stronger population growth than cities within 100 to 200 km of container ports.

In order to address the concern that population growth is a function of initial size, Column 4 adds a control for 1950 log population. This specification allows comparisons of population growth between cities of similar initial sizes. In this specification, cities 0 to 100 km from a container port experience a statistically significant 14.54 percentage points greater population growth, which is about 9 percent of the average city growth over the

---

[15] We plan to provide statistics for this claim in future drafts.

period. Comparing these results with Column 3, we interpret this change as evidence that initially large cities grow more slowly, and that growth due to containerization is concentrated in larger cities. Also, in this final column, population growth associated with containerization is concentrated in cities that are nearest (less than 100 km from) to containerized ports. The remaining two coefficients are both small and imprecisely estimated.

While the OLS specification does mitigate many possible endogeneity concerns, the possibility remains that a time-varying factor causes cities to be both near a container port and to have higher population growth. To address this concern, we turn to instrumental variable estimates.

We begin with Table 3a, which presents the full first stage estimates. For the world sample, we instrument the three proximity measures with three measures of proximity to a port of depth greater than 30 feet, using the same distance bins as the containerization proximity measures.

If the instrument works as we hypothesize, we expect strong and significant results along the diagonal of the table – that is, proximity within 100 km to a very deep port should be strongly correlated with proximity within 100 km of a container port. We see this hypothesized pattern very strongly in each of the three regressions in the first panel of the table, where the dependent variable is proximity to a containerized port by distance. In the first column, the estimates report that cities within 100 km of a deep port in 1953 are 55 percent more likely to be within 100 km of a container port. This increase is relative to a mean of 35 percent of cities being within 100 km of a container port—a sizeable increase.

The estimates in the other two equations are of roughly similar magnitudes and significance. The off-diagonal coefficients in this table are generally negative. This suggests that there is some geographic competition in the location of container ports. Intuitively, we would expect that increasing the number of suitable locations to build a containerized port would decrease the likelihood of containerization of every location. The negative coefficients off-diagonal are consistent with this. In all columns, the F statistic is quite high, and is never less than 62.

In the right panel of the table, we repeat the same estimations, using a city's years since first containerization by distance bin as the dependent variable. Again, results on the diagonal are strong and positive. Cities within 100 km of a very deep 1953 port have, on average, 21 years more exposure to a container port within 100 km. This is a large estimate, given that the mean years to first containerization is 11. In this second panel, the F statistic is never lower than 64.

In sum, this table suggests that the first condition for instrument validity is satisfied: there is a strong relationship between the instrument and the endogenous variable. We are limited by the extent of the world data in our ability to test the second condition for instrument validity—that proximity to a very deep port impacts population growth only through proximity to container ports. We return to further tests of this second condition with the U.S. sample.

Moving to the two-stage least squares results in Table 4a, the results show a very similar pattern to the OLS panel as we add covariates. In general, the OLS results are somewhat larger than the IV results. In our most complete specification (column 8), cities within 100 km of a container port show a precisely estimated increase in population growth of about 25 percentage points, or about 16 percent of the mean. For cities within 100 to 200 km and within 200 to 300 km of a container port, we estimate a population growth increase of 23 percentage points.

Why are the IV results larger than the OLS results? Suppose that larger cities grow at a slower rate than smaller cities, and larger cities are more likely to be near containerization adopting ports. When we correct for this endogeneity with the instrument—in principle, giving larger weight to smaller cities, where the depth is the main driver of the containerization decision—the coefficient should increase.

The specification in Table 4a measures proximity to containerized ports with a dummy variable. While this has the benefit of being easy to interpret, it reports an average across cities near a container port for very few years and cities near a container port for many years. If cities near early-adopting container ports have different population growth trajectories than cities near later-adopting container ports, the average results could be quite misleading. For example, if there are many later-adopting container ports, the average effect could be quite small, even while some cities experience large effects. Alternatively, if proximity to containerization requires a certain number of years to achieve population growth, or stops after a number of years, the average could again be misleading.

We address this issue by estimating the impact of each additional year since first containerization, measured by proximity to a container port. Table 5a reports estimates from this specification. Results here have a similar pattern of sign and significance as in Table 4a. The final OLS column (column 4) estimates that cities within 100 km of a container port grow, on average, ($0.004 * 32.25$, where 32.25 is the average years since first containerization in this distance bin) 12.9 percentage points faster than cities never near a container port, and 8 percent faster than the average city. This is similar to the OLS estimate of a 14.5 percentage point increase from Table 4a.

The final IV column (column 8) suggests that cities within 100 km and 100 to 200 km of a container port grow (0.0067*32.25) 21.6 percentage points and (0.0063*30.5) 19.2 percentage points faster. These imply growth 14 percent faster than the average city. Again, these estimates are a little smaller than the results using "ever containerizing" as the key measure. This suggests that cities that have been exposed to container ports longer have smaller (per year) impacts. If larger cities are more likely to be near ports that adopt container technology early, and larger cities grow more slowly, this could explain the decline. It is also possible that, over time, there are decreasing marginal returns from proximity to a container port, so that beyond a certain point additional years do not contribute to additional population growth.

We conclude our analysis of the world sample by investigating potential concerns with our identification. One concern is our use of all cities for counterfactual population growth. Perhaps port cities—specifically, cities near 1953 ports—grow differently than all cities, conceivably because they are all more affected by industries that specialize in waterborne trade. Under this assumption, other port cities could be the better counterfactual. (Nonetheless, one could equally argue that the set of all cities provides a better counterfactual, particularly in countries with one very large port city for which the other port cities are a poor counterfactual.) To test this contention, we re-estimate the results using only cities within 300 km of a 1953 port. The first column in Table 6a repeats the most complete instrumental variable specification from Table 4a for comparison. The following column limits the sample to only port cities, where we find very similar results. Indeed, in each column pair in this table, we use the full sample of cities followed by the sample of port cities only (more on the remaining columns in the following paragraph). Regardless of specification, the results in the paired columns are extremely similar, and are not qualitatively differentiable from one another.

An additional concern with the world sample is that the rule for entry into the sample—a population greater than 300,000 at any point from 1950 to 2014—biases the sample toward faster growing cities, and could lead us to overestimate containerization's impact on urban growth. We address this concern by limiting the sample to cities that have 300,000 people or more in all years of the sample. This is, in effect, a sample without any selection biases due to growth rates. It is a more limited sample of cities, but one without interpretation concerns.

Regardless of whether we use all cities in this range (303 cities, column 3) or only port cities (213 cities, column 4), we still find that cities near ports have faster population growth after the advent of containerization. However, for the larger cities, we see that growth is primarily associated with being only somewhat close to containerized ports. For these larger cities, we estimate that being 100 to 200 km from a port causes roughly 41 percentage points more growth, or a 40 percent increase relative to the mean. We speculate that this change in distance pattern may be due to initially larger cities having less room to accommodate the requirements of a container port—which are substantial in terms of land area—in their immediate vicinity.

The final two columns of this table repeat the specification using a city's years of proximity to a container port by distance bin as the key endogenous variable. Here, again, we see the patterns of both port cities and the full sample yielding very similar results. As in the first panel, limiting the sample to only larger cities (columns 7 and 8) yields larger estimates—a 34.2 percentage point increase in population growth, or 33 percent relative to the mean (column 7). Moreover, cities slightly farther from container ports have a statistically significant increase in population growth.

5.2 U.S. Results

Having explored containerization's impact on large world cities, we now turn to the U.S. sample, where we can assess containerization's impact on areas of all sizes. The righthand panel of Table 4b reports the OLS coefficients for the estimation of equation 1. The first column controls only for the presence of a port in 1953 within 300 kilometers. Results from this specification show that counties nearest to containerized ports have the largest absolute increases in population growth. From 1950 to 2010, their population growth was 31 percentage points higher than in counties never near a container port. This is 30 percent of the average change in population for all counties. Column 2 adds a vector of controls for initial port intensity: the number of 1953 ports in each of the six distance bins. This inclusion increases the coefficient for counties within 50 kilometers of container ports, but leaves the other coefficients relatively unchanged, suggesting that containerization adoption is not strongly related to the presence of many nearby ports in 1953.

To address the issue that cities of different sizes may grow at different rates—and that cities near container ports are larger, as we know from Table 1b—the third column adds controls for log population in 1920, 1930, 1940 and 1950. In other words, we allow differential growth rates by initial city size. Here initial size is not just in the first pre-treatment year, but an additional thirty years preceding the treatment. The addition of these controls decreases the estimates of container port proximity on population growth by somewhat less than 50 percent for counties very close (less than 100 km) to container ports, and has a smaller effect on counties farther from container ports.

Finally, as we saw in Table 1b, counties near future container ports had, on average, much higher rates of manufacturing employment in 1956. The final column includes this variable as a control. It also attempts to control for pre-containerization port prominence by including, for all ports in each distance bin, total 1955 international trade in millions of dollars. The addition of these covariates has little effect on the coefficients. In this final specification, counties within 50 km of a port that containerizes experience a statistically significant 24 percentage point increase in population growth, which is about one-quarter of the average change in the dependent variable. Cities within 50 to 100 km of a container port experience a statistically significant additional 15 percentage point growth, or a 16 percent increase relative to the mean. Cities 100 to 150 km from a container port experience 13 percent greater growth, and cities 150 to 200 km from a container port experience additional relative growth of 12 percent.

As we remain concerned that an additional, time varying factor may cause both proximity to containerization and population growth, we turn to our instrumental variable strategy. Table 3b reports estimates from the first stage, using the maximal set of controls from Table 4b. The U.S. specification has six endogenous variables—proximity to container ports at the six distance bins—and six instruments, which are the depth of the deepest port in 1953 in each of the six distance bins. As with the world table, we expect that the relationship in this table should be strongest on the diagonal: proximity to a container port at distance d_1 to d_2 should be most correlated with the 1953 depth of the deepest port in that same distance interval.

This is in fact the pattern we see. In panel A, where the dependent variable is proximity to an ever-containerized port, counties in proximity to very deep ports are between 33 and 43 percent more likely to be near a container port at the same distance. The coefficient is smallest at the 250 to 300 distance. All coefficients on the diagonal are strongly significant, and the F statistics for these regressions are never lower than 44; all but five are about 100.

The bottom panel of the table uses years of proximity to the first container port by distance bin as the dependent variable. In this specification, counties located within 0 to 50 km of a very deep 1953 port experience an additional 26 years of proximity to a container port. The instrument explains roughly half of the 41 average years of proximity in this distance bin. Coefficients on the other bins are slightly larger, in the 28 to 32 year range, save the final coefficient for the 250 to 300 km bin, which is 19 years. All coefficients are significant at the 1 percent level or above, and the F statistics for each estimation are never lower than 92.

The U.S. data allow us to further test whether the instrument is valid. To do so, we evaluate whether the instruments are correlated with county-level characteristics that might plausibly be in the error term. While we cannot do this for all potential confounders, we can observe whether the identifying variation—the residual from a regression of the instrument on the full set of covariates—is correlated with specific pre-treatment covariates. Figure 3 uses the instrument $\mathbb{1}\{\text{Very deep port in 1953 between 0 and 50 km}\}_i$.

Our regression specification controls for log of population density in 1920, 1930, 1940 and 1950. Were the identifying variation in the instrument to be related to the log of 1910 population density, this would suggest that the pre-treatment controls were not adequately capturing the historical pattern of population growth. We do not find this to be the case. Figure 3a shows the identifying variation from the instrument on the y axis, and log of 1910 population density on the horizonal axis. We find no significant relationship ($t = 0.55$) between these two variables.

Similarly, recall that the regression controls for the 1955 value of international trade at a set of distances from each county. If this covariate did not sufficiently control for the impact of pre-containerization port strength on population, we would expect that the identifying variation would be related to the 1948 value of international trade at a set of distances from each county. Figure 3b shows, this is not the case. Again, the relationship between the identifying variation and the variable of concern is insignificant ($t = -0.95$).

In fact, we do twelve estimations: a regression of the identifying variation from each of the six distance bins with 1910 population, and with the dollar value of 1948 international trade at ports in that distance interval. In these 12 regressions, we find one significant relationship. This one significant coefficient is almost what we would expect only by random chance.

Having allayed concerns about instrument strength and validity, we turn to the second half of Table 4b, which shows results from the instrumental variable estimation. As with the world sample, the pattern of coefficient change as we add covariates is very similar between the IV and OLS estimates. One notable difference in the U.S. estimation is that the IV results show significant effects at the 250 to 300 km distance range, while the OLS results do not. Counties at this distance from a port have statistically significantly more population growth. In the final column, counties closest to ports (0 to 50 km) have 40 percentage points (and 40 percent) more population growth. Counties 100 to 150 km of a port see no statistically significant difference in growth and then cities 150 to 200 km see an additional 17 percentage points of population growth, or 27 percent of the average change over the period. Cities located from 200 to 250 kilometers of a port see an additional 28 percentage points of population growth, or a 49 percent increase relative to the mean.

The U-shaped pattern of the coefficients with respect to distance from treatment is a somewhat surprising feature of the U.S. results. From a standard new economic geography model (e.g. Redding, 2015), we would expect the effect of containerization to be largest for counties closer to containerized ports. However, the non-linear geographic effect of treatment may be explained by the intermodal nature of containers. In the United States, more than anywhere else in the world, containerization impacted the entire transit network. As an intermodal system, firms can reap the benefits of better access to international markets without necessarily being very close to a container port. As such, the counties that gained the most from the new shipping technology are those near the eight percent of U.S. ports that containerize, as well as counties farther from container ports with more to gain from the new shipping technology in terms of access to international markets. That the treatment effect follows a non-linear pattern could also reflect spillovers across space if people migrate to grab some of the rent associated with large infrastructure investments concentrated near container terminals. To the extent that people are more likely to migrate to nearby destinations, we would expect treatment effects to be non-linear.

As with the world sample, the IV results here are larger than the OLS ones. The cause for the divergence may also be similar. If already populous counties grow at a slower rate than less populous ones, and the instrument gives more weight to smaller cities, where proximity to containerization is more likely to be driven by the supply constraint posed by depth, this is what we should expect.

Estimates of proximity to a containerized port on population growth, using the dichotomous specification in Table 4b, might produce misleading results if population increases are limited to counties near early-adopting (or later-adopting) ports. To address this concern, Table 5b reports results from re-estimating the specifications in Table 4b, using years of proximity to a containerized port by distance bin as the dependent variable.

Concentrating on column 8, our most complete instrumental variables estimate, being a county 0 to 50 km from a container port is associated with a (0.007*41) 29 percentage point increase in population growth, which for this distance category is also a roughly 30 percent increase relative to the mean. Counties 150 to 200 km experience (0.003*39/0.63) 19 percent more population growth, and counties 200 to 250 km experience an additional (0.005*40/0.57) 35 percent growth, relative to the mean in each distance bin.

Comparing these results to those using the proximity measured dichotomously, magnitudes are generally somewhat smaller. We hypothesize that the same factors that drove the world sample estimates using years of proximity to a container port to be smaller are also at play here.

We use Table 6a to explore the robustness of our instrumental variable estimates to alternative sampling frames. As with the world sample, one might contend that the proper counterfactual counties are those near 1953 ports. Column 1 limits the sample to these 1,296 counties only. Estimates in this specification are generally smaller and less significant than those in Table 4b's column 8. This suggests that U.S. counties near 1953 ports grow faster, on average, than counties far from 1953 ports. Given the high correlation between proximity to a 1953 port and proximity to a containerized port, a comparison of columns 7 and 8 in the summary statistics table shows this to be the case.

Despite the decline in magnitude, the pattern of larger growth changes when the county is either quite close to a container port (0 to 50 km) or not particularly close (200 to 250 km) holds.

To explore how much our results are influenced by the correlation between county population and proximity to container ports, we split the sample into counties that in 1950 were above and below median population. As the sample size declines, we lose the ability to obtain precise estimates at many distance intervals. However, a comparison of the coefficients for counties closest to ports (0 to 50 km, comparing columns 2 and 3) shows that counties that had below median population in 1950 grew much more quickly after the advent of containerization. In counties with 1950 population above the median, growth is primarily associated with being not particularly proximate (200 to 250 km) to a container port. As in the world sample, it may be that these more populous counties cannot, or are too expensive, to house a container port in the direct vicinity, but still benefit from the trade deriving from the port.

The final panel of the table repeats these estimates using years of proximity to a container port as the endogenous variable. We find a very similar geographic pattern to the specification using the proximity to an ever containerized port, and the coefficients yield changes of very similar magnitudes. This suggests that the ever-containerized results are not distorted by unusual patterns of population growth in response to additional years of containerization.

5.3 Comparison of World and U.S. Results

The world and U.S. results both show that proximity to a container port, at certain distances, is associated with population growth above the average. Estimates for the world sample are somewhat smaller in magnitude than for the U.S. sample, and differ in their geographic pattern. There are multiple possible reasons for larger results in the U.S. First, the U.S. was the first country in the world to adopt container shipping; the average U.S. county proximate to a container port has experienced ten more years since first containerization than the average world city. In addition, U.S. container trade was initially primarily domestic, which may not be true of other countries, and which may yield additional population growth impacts. Finally, population in U.S. counties is substantially smaller than the average population in the world cities data. If smaller cities grow more rapidly, this could account for part of the difference.

The geographic pattern of city proximity to container ports and urban growth also differs in the world and U.S. samples. This comparison is difficult to make with precision, however, given the difference distance bands we have used in the two analyses. We defer analysis of this difference for a future draft.

6 Conclusion

We use U.S. and world data from 1950 to 2010 to assess the long-term consequences of containerization. This technology not only transformed global trade, but had substantial local consequences. We find substantial increases in population for cities closest to container ports, and still sizeable increases for cities at a middle distance from container ports.

In future work, we would like to push the reduced form methods to probe the drivers of population change, and containerization's impact on other local variables. However, the reduced form methods do not allow us to assess whether population increases near ports comes at the expense of other locations. To tackle this general equilibrium proposition, we plan to turn to a market access analysis, as in Donaldson and Hornbeck (forthcoming). We are currently assembling historical maps for this purpose, and anticipate having the results in a future draft.

References

Autor, David H., Dorn, David, and Hanson, Gordon H., 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103(6): 2121–68.

Barro, Robert J, 1991. "Economic Growth in a Cross Section of Countries." *The Quarterly Journal of Economics* 106(2): 407–43.

Bernhofen, Daniel, El-Sahli, Zouheir, and Kneller, Richard, forthcoming. "Estimating the Effects of the Container Revolution on World Trade." *Journal of International Economics* .

Bridgman, Benjamin, 2014. "Why Containerization Did Not Reduce Ocean Trade Shipping Costs." Working Paper.

Broeze, Frank, 2002. *The Globalisation of the Oceans: Containerisation From the 1950s to the Present*, vol. 23 of *Research in Maritime History*. St. John's, Newfoundland: International Maritime Economic History Association.

Davis, Donald and Weinstein, David, 2002. "Bones, Bombs, and Break Points: The Geography of Economic Activity." *American Economic Review* 92(5): 1269–1289.

Dix-Carneiro, Rafael and Kovak, Brian, 2015. "Trade Reform and Regional Dynamics: Evidence From 25 Years of Brazilian Matched Employer-Employee Data." Working paper.

Donaldson, Dave, forthcoming. "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." *American Economic Review* .

Donaldson, Dave and Hornbeck, Richard, forthcoming. "Railroads and American Economic Growth: A Market Access Approach." *Quarterly Journal of Economics* .

Duranton, Gilles, Morrow, Peter M., and Turner, Matthew A., 2014. "Roads and Trade: Evidence from the US." *Review of Economic Studies* 81: 681–724.

Duranton, Gilles and Turner, Matthew, 2012. "Urban Growth and Transportation." *Review of Economic Studies* 79: 1407–1440.

Feyrer, James, 2009. "Trade and Income – Exploiting Time Series in Geography." NBER Working Papers 14910, National Bureau of Economic Research, Inc.

Fujita, Masahisa, Krugman, Paul, and Venables, Anthony J., 1999. *The Spatial Economy: "Cities, Regions and International Trade"*. The MIT Press.

Gonzalez-Navarro, Marco and Turner, Matthew, 2015. "Subways and Urban Growth: Evidence from Earth." Working paper.

Hanson, Gordon, 2005. "Market Potential, Increasing Returns, and Geographic Concentration." *Journal of International Economics* 67: 1–24.

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

Hummels, David and Schaur, Georg, 2013. "Time as a Trade Barrier." *American Economic Review* 103(7): 2935–2959.

Kendall, L.C., 1986. *The Business of Shipping*. Cornell Maritime Press, fifth edn.

Kondo, Illenin O., 2013. "Trade Reforms, Foreign Competition, and Labor Market Adjustments in the U.S." International Finance Discussion Papers 1095, Board of Governors of the Federal Reserve System (U.S.).

Krugman, Paul, 1991. "Increasing Returns and Economic Geography." *Journal of Political Economy* 99(3): 483–99.

Levinson, Mark, 2008. *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger*. Princeton, NJ: Princeton University Press.

Lucas, Robert Jr., 1988. "On the mechanics of economic development." *Journal of Monetary Economics* 22(1): 3–42.

Michaels, Guy, 2008. "The effect of trade on the demand for skill: Evidence from the Interstate Highway System." *Review of Economics and Statistics* 09: 683–701.

Morgan, F. W., 1952. *Ports and Harbours*. London, UK: Hutchison's University Library.

Morison, Robert F., 1977. "Containerization: Good and Bad News." *The Journal of Commerce and Commercial* pages 1A, 15A.

National Geospatial Intelligence Agency, 1953. "World Port Index, 1953." H.o. pub. no. 950, first edition.

National Geospatial Intelligence Agency, 2015. "World Port Index, 2015." Twenty-fourth edition.

Pascali, Luigi, 2014. "The Wind of Change: Maritime Technology, Trade and Economic Development." Working Paper 165, CAGE.

Ramondo, Natalia, Rodriguez-Clare, Andres, and Saborio-Rodriguez, Milagro, 2014. "Trade, Domestic Frictions, and Scale Effects." Working paper.

Redding, Stephen and Sturm, Daniel M., 2008. "The Costs of Remoteness: Evidence from German Division and Reunification." *American Economic Review* 98(5): 1766–1797.

Redding, Stephen J., 2015. "Goods Trade, Factor Mobility and Welfare." NBER Working Papers 18008, National Bureau of Economic Research, Inc.

Redding, Stephen J. and Turner, Matthew A., 2015. "Transportation Costs and the Spatial Organization of Economic Activity." In Gilles Duranton and William Strange, (Eds.) "Handbook of Regional and Urban Economics," Elsevier B. V.

Rodrigue, Jean-Paul, 2015. `https://people.hofstra.edu/geotrans/eng/ch3en/conc3en/USintermodalrail.html`.

Romer, David H. and Frankel, Jeffrey A., 1999. "Does Trade Cause Growth?" *American Economic Review* 89(3): 379–399.

Rua, Gisela, 2014. "Diffusion of Containerization." Finance and Economics Discussion Series 88, Board of Governors of the Federal Reserve System (U.S.).

Sargent, A. J., 1938. *Seaports and Hinterlands*. London, UK: Adam and Charles Black.

Talley, Wayne K., 2002. "Dockworker Earnings, Containerisation, and Shipping Deregulation." *Journal of Transport Economics and Policy* 36(3): 447–467.

Topalova, Petia, 2010. "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India." *American Economic Journal: Applied Economics* 2(4): 1–41.

United Nations Conference on Trade and Development, 2013. *Review of Maritime Transport*. New York, NY: United Nations.

Wilson, R. P., 1982. "The Containerization of Ports." UC Santa Cruz Working Paper.

Figure 1: Adoption of Containerization: 1956–2008

(a) Worldwide

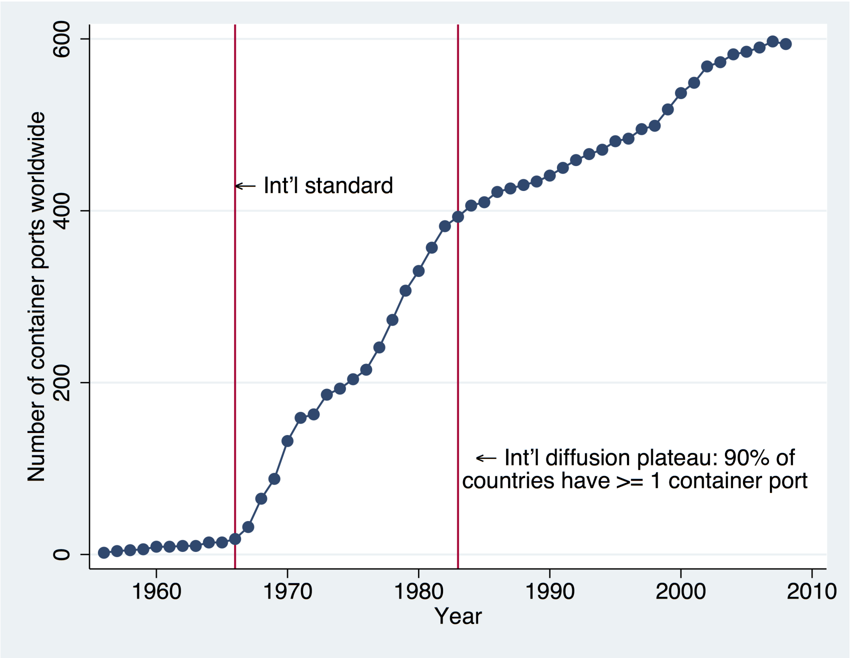

(b) United States

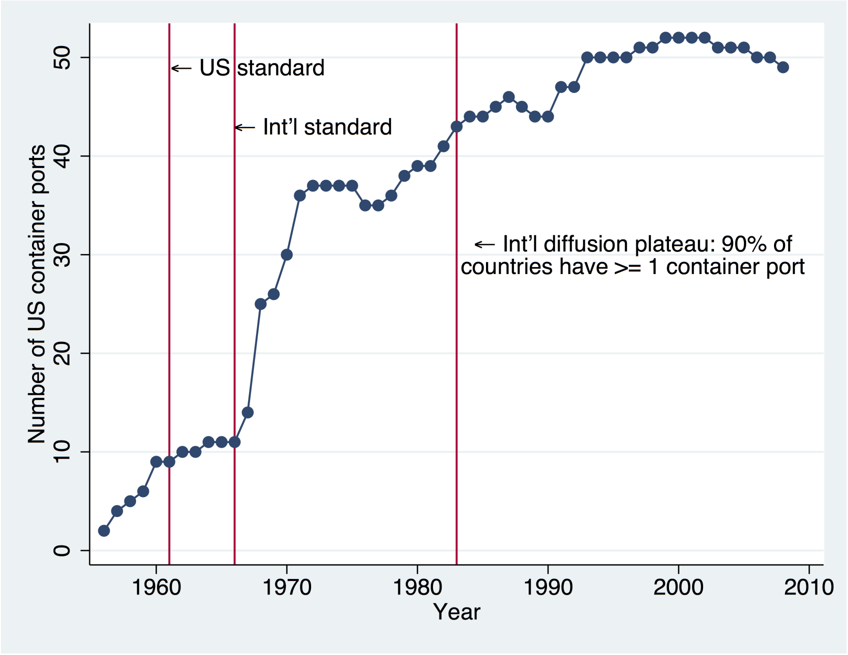

Source: *Containerisation International Yearbook*, volumes 1968 and 1970–2010.

Figure 2: Likelihood of Having a Containerized Port by 1953 Port Depth

(a) Worldwide

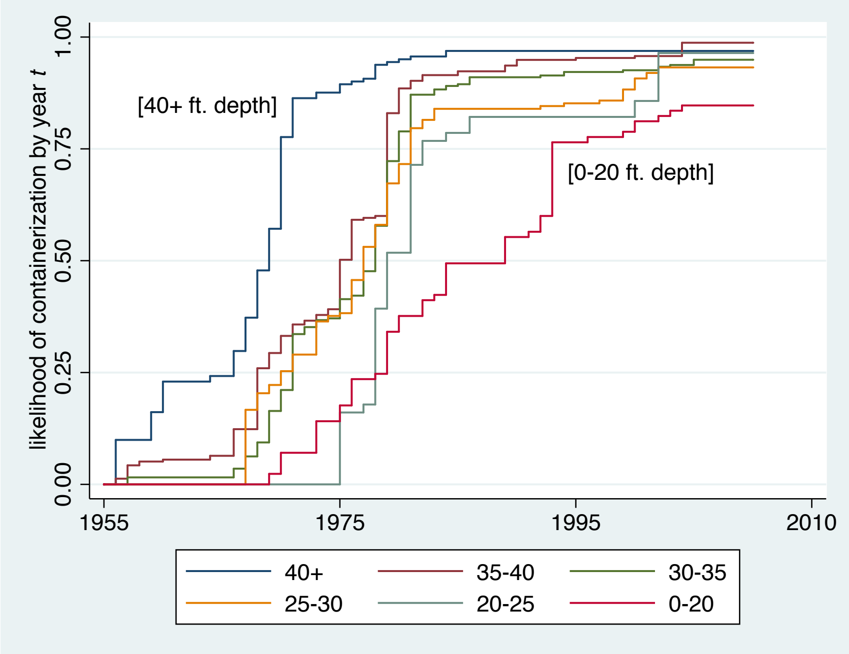

(b) United States

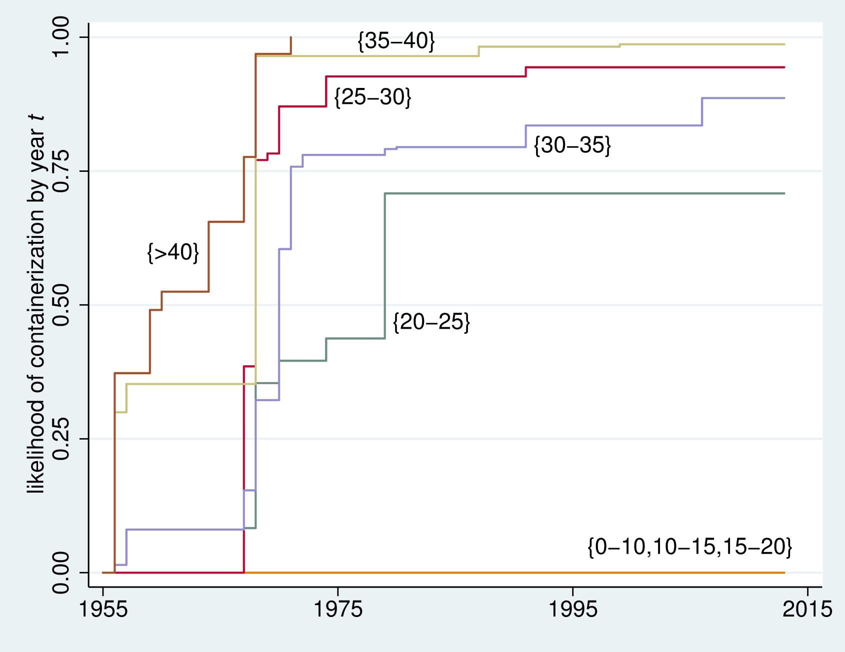

Notes: In this picture, a city (world) or county (U.S.) has a port if there is a port within 300 km. We call this city or county "containerized" if $t >$ year of first containerization of any port within 300 km. We measure depth as the depth of the deepest port within 300 km. On average, deeper ports are more likely to ever containerize, and more likely to containerize early.

Figure 3: Instrument Variation vs. Pre-Treatment Covariates

(a) Versus Log of Population Density, 1910

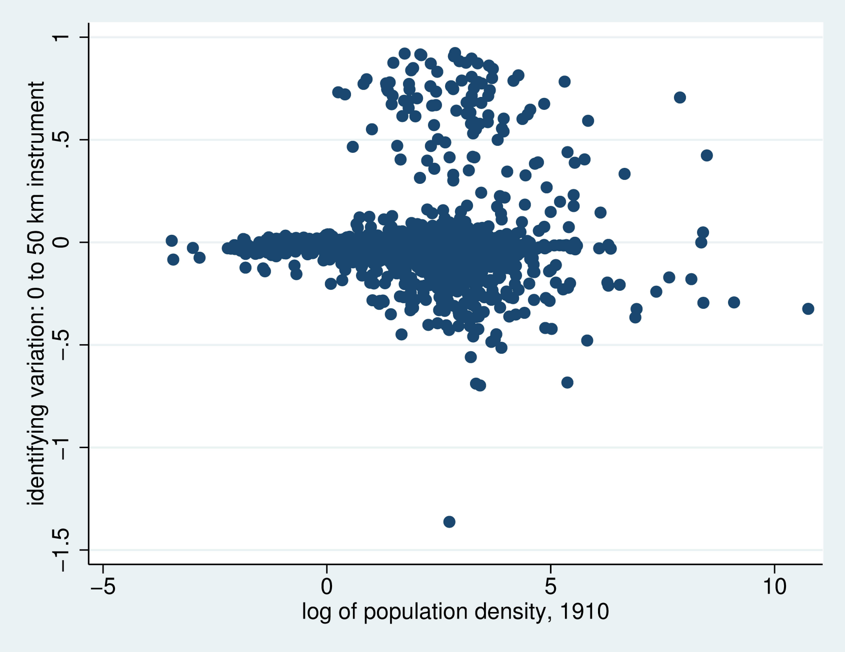

(b) Versus Millions of Dollars of 1948 International Trade at Port within 50 to 100 km

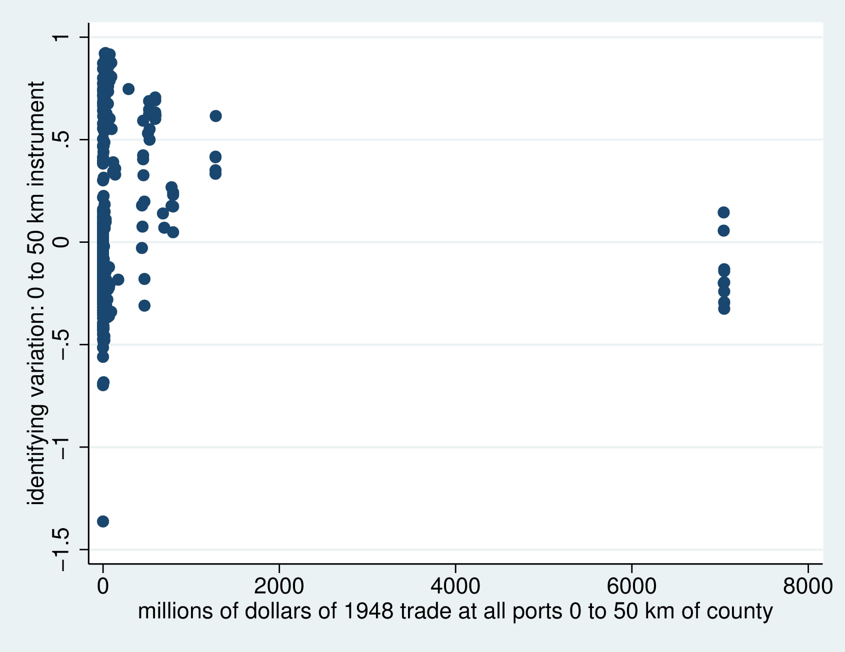

Notes: “Identifying variation” is the residual from a regression of the instrument (county is within 0 to 50 kilometers of a “very deep” port) on the full set of covariates from equation 1 (as in Table 4, Columns 4 and 8).

Figure 4: Containerization's Relationship with Population Strongest at Closer Distances

(a) Worldwide

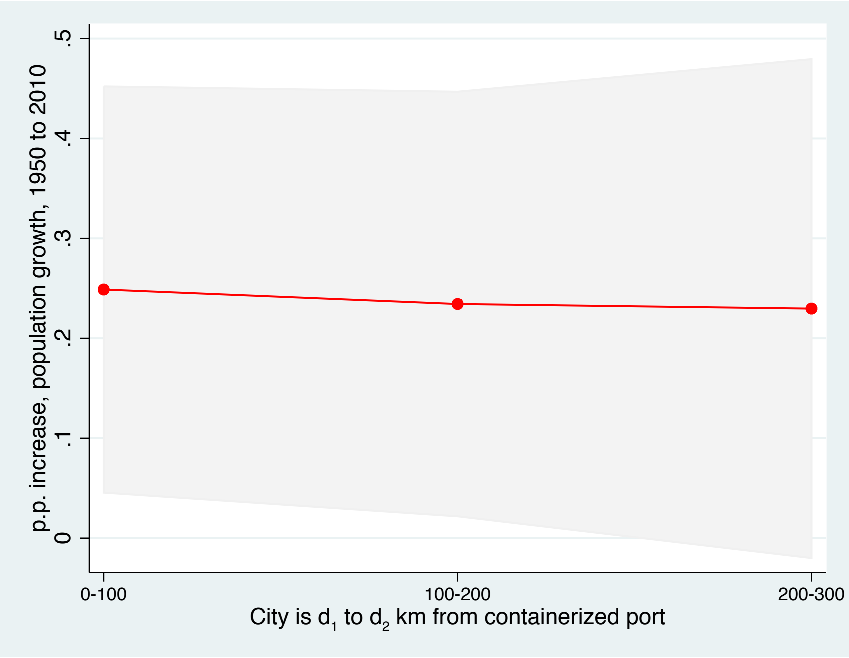

(b) United States

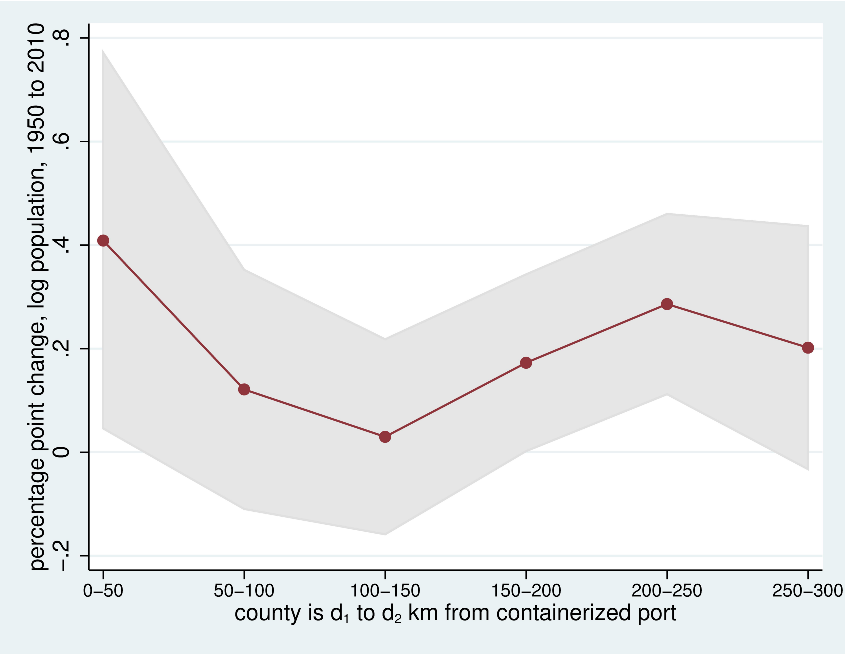

Notes: This picture presents results from Column 8 of Table 4; each dot corresponds to an estimated coefficient for each distance bin. The gray band is the 95% confidence interval.

Table 1: City Characteristics by Distance to Containerized Port

(a) Worldwide

	Distance to Containerized Port							
	0 to 50	50 to 100	100 to 150	150 to 200	200 to 250	250 to 300	Ever Cont.	Never Cont.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Population								
1950	12.674	12.493	12.399	12.365	12.331	12.321	12.322	11.985
	[1.147]	[1.087]	[1.105]	[1.081]	[1.003]	[1.028]	[1.058]	[0.811]
2010	14.112	13.799	13.722	13.666	13.653	13.666	13.810	13.603
	[1.067]	[1.026]	[0.999]	[0.962]	[0.903]	[0.917]	[0.978]	[0.804]
Continent								
Africa	0.101	0.055	0.056	0.100	0.055	0.067	0.097	0.048
Asia	0.348	0.346	0.372	0.363	0.376	0.412	0.378	0.585
Australia	0.033	0.011	0.017	0.012	0.000	0.011	0.014	0.000
Europe	0.272	0.313	0.325	0.327	0.345	0.326	0.241	0.186
North America	0.156	0.203	0.175	0.143	0.188	0.139	0.184	0.115
South America	0.091	0.071	0.056	0.056	0.035	0.045	0.087	0.067
Years Since	32.438	31.852	30.774	30.378	31.176	30.899	35.171	.
First Cont.	[11.158]	[11.790]	[12.406]	[11.845]	[12.125]	[11.513]	[9.667]	.
Observations	276	182	234	251	255	267	632	419

Notes: The unit of observation in this table is the city. We report means and standard deviations (brackets) for each variables.
Source: See data appendix.

Table 1: City Characteristics by Distance to Containerized Port

(b) United States

	Distance to Containerized Port							
	0 to 50	50 to 100	100 to 150	150 to 200	200 to 250	250 to 300	Ever Cont.	Never Cont.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Population								
1910	10.85	10.23	10.2	10.14	10.11	10.11	10.15	9.5
	[1.45]	[1.09]	[1.06]	[0.95]	[0.86]	[0.91]	[0.95]	[0.89]
1920	11.05	10.34	10.29	10.21	10.16	10.16	10.21	9.56
	[1.49]	[1.15]	[1.1]	[1]	[0.92]	[0.96]	[1]	[0.92]
1930	11.29	10.48	10.4	10.29	10.22	10.21	10.28	9.64
	[1.52]	[1.23]	[1.16]	[1.06]	[0.98]	[1.02]	[1.06]	[0.84]
1940	11.41	10.59	10.48	10.37	10.29	10.28	10.36	9.68
	[1.5]	[1.24]	[1.17]	[1.07]	[0.98]	[1.03]	[1.07]	[0.86]
1950	11.67	10.76	10.61	10.46	10.36	10.34	10.44	9.66
	[1.52]	[1.33]	[1.23]	[1.14]	[1.06]	[1.1]	[1.15]	[0.92]
2010	12.63	11.66	11.33	11.09	10.93	10.85	11	9.89
	[1.3]	[1.45]	[1.36]	[1.32]	[1.28]	[1.28]	[1.34]	[1.29]
Region								
Northeast	0.25	0.18	0.2	0.18	0.16	0.16	0.16	0.01
Midwest	0.18	0.2	0.22	0.26	0.27	0.28	0.31	0.41
South	0.41	0.45	0.45	0.45	0.49	0.49	0.44	0.44
West	0.17	0.18	0.13	0.11	0.09	0.07	0.09	0.15
Share Manuf.	0.45	0.42	0.43	0.43	0.43	0.43	0.43	0.27
Emp., 1956	[0.17]	[0.18]	[0.18]	[0.18]	[0.18]	[0.18]	[0.18]	[0.22]
Years Since	41.24	39.69	40.13	39.37	39.81	40.34	44.36	.
First Cont.	[9.19]	[10.1]	[9.59]	[9.85]	[9.19]	[9.11]	[7.6]	.
Observations	120	256	392	522	614	692	1209	1493

Notes: The unit of observation in this table is the county. We report means and standard deviations (brackets) for each variables.
Source: See data appendix.

Table 2: Containerization More Likely in Cities or Counties Near Deeper Ports

(a) Worldwide

	Ever Cont.		Years Since Cont.	
	(1)	(2)	(3)	(4)
40+ Depth	0.524**		18.022**	
	(0.205)		(8.117)	
35-40 ft. Depth	0.533***		16.414**	
	(0.205)		(8.077)	
30-35 ft. Depth	0.515**		15.754*	
	(0.206)		(8.098)	
25-30 ft. Depth	0.443**		9.878	
	(0.208)		(8.176)	
20-25 ft. Depth	0.486**		2.389	
	(0.211)		(8.392)	
15-20 ft. Depth	0.154		-2.885	
	(0.252)		(8.662)	
10-15 ft. Depth	0.413*		0.898	
	(0.214)		(8.260)	
30+ ft. Depth		0.105***		9.425***
		(0.033)		(1.269)
Mean, Dependent Variable	0.60	0.60	21.15	21.15
R-squared	0.89	0.89	0.88	0.87
F Stat Excluded Instrument(s)	3.14	10.00	18.17	55.12
Observations	1051	1051	1051	1051

Sample: World cities over 50K in 1950.
Dependent variable (1)-(2): Adoption of containerization within 300km.
Dependent variable (3)-(4): Number of years since the adoption of containerization within 300km.
Port depth is the depth of the deepest port within 300km.
All specifications control for country fixed effects, log population in 1950,
a dummy variable for having a port within 300km in 1953, and the number
of ports within 300km in 1953.
Robust standard errors in parentheses. Stars denote significance levels: * 0.10 ** 0.05 *** 0.01.

Table 2: Containerization More Likely in Cities or Counties Near Deeper Ports

(b) United States

	Dependent Variable is			
	Ever Containerized		Years Since First Cont.	
	(1)	(2)	(3)	(4)
Port Depth in feet is				
40 and over	1.028^{***}		46.435^{***}	
	(0.064)		(3.119)	
35-40	0.999^{***}		45.434^{***}	
	(0.063)		(3.08)	
30-35	0.888^{***}		35.199^{***}	
	(0.062)		(3.067)	
25-30	0.959^{***}		40.709^{***}	
	(0.063)		(3.072)	
20-25	0.705^{***}		27.323^{***}	
	(0.065)		(3.214)	
15-20	-0.012		0.183	
	(0.107)		(5.253)	
10-15	0.002		0.103	
	(0.084)		(4.127)	
$1\{\text{Depth} \geq 30 \text{ Feet}\}$		0.069^{***}		3.793^{***}
		(0.011)		(0.532)
Mean, Dependent Variable	0.447	0.447	19.85	19.85
R-squared	0.909	0.884	0.894	0.864
Observations	2702	2702	2702	2702
F for Excluded Instrument(s)	111	41	116	51
Increase in R^2 due to instrument	0.026	0.002	0.032	0.003

Notes: We report standard errors below coefficients in parentheses. + indicates significance at the 10% level, * at the 5% level, ** at the 1% level, and *** at the 0.1 % level. Port depth is the depth of the deepest port within 300 km. All specifications control for a dummy for ever being within 300 km of a 1953 port, the number of 1953 ports in each of six distance bins to 300 km, log population 1920 to 1950, 1956 manufacturing share, and the total value of waterborne international trade in each of six distance bins to 300 km. This table reports results using wharf depth; we show similar results for channel depth and anchorage depth in Appendix Table 4.
Source: See data appendix.

Table 3: First Stage: Containerization More Likely When Ports are Deep

(a) Worldwide

	Ever Cont.			Years since Cont.		
	(1) 0 to 100	(2) 100 to 200	(3) 200 to 300	(4) 0 to 100	(5) 100 to 200	(6) 200 to 300
30+ ft. Depth (0-100km)	0.552***	-0.095**	-0.045	21.236***	-3.558***	-2.925**
	(0.036)	(0.038)	(0.040)	(1.413)	(1.328)	(1.397)
30+ ft. Depth (100-200km)	-0.086**	0.499***	-0.066	-2.895**	18.119***	-2.820**
	(0.037)	(0.036)	(0.040)	(1.338)	(1.425)	(1.366)
30+ ft. Depth (200-300km)	-0.106***	-0.096**	0.484***	-2.817*	-2.457*	19.454***
	(0.041)	(0.042)	(0.040)	(1.456)	(1.463)	(1.498)
Mean, Dependent Variable	0.35	0.35	0.37	11.71	11.08	12.02
R-squared	0.65	0.63	0.63	0.68	0.66	0.67
F Stat Excluded Instrument(s)	99.72	73.97	62.39	97.77	64.31	84.79

Sample: World cities over 50K in 1950.
Dependent variable (1)-(3): Adoption of containerization between d1 and d2 km.
Dependent variable (4)-(6): Number of years since the adoption of containerization between d1 and d2 km.
Port depth is the depth of the deepest port within 300km.
All specifications control for country fixed effects, log population in 1950,
a dummy variable for having a port within 300km in 1953, and the number
of ports within 300km in 1953.
All regressions have 1051 observations.
Robust standard errors in parentheses. Stars denote significance levels: * 0.10 ** 0.05 *** 0.01.

Table 3: First Stage: Containerization More Likely When Ports are Deep

(b) United States

	Dependent Variable: County is d_1 to d_2 of containerized port					
	0 to 50	50 to 100	100 to 150	150 to 200	200 to 250	250 to 300
	(1)	(2)	(3)	(4)	(5)	(6)
A. Dependent Variable is Ever Containerized						
Depth is $\geq$ 30 feet, county is d_1 to d_2 km of any port, $t > 1956$						
0 to 50	0.427***	0.029	-0.023	-0.034	-0.103**	-0.008
	(0.019)	(0.026)	(0.031)	(0.035)	(0.037)	(0.039)
50 to 100	0.030*	0.521***	0.021	-0.028	0.039	-0.018
	(0.015)	(0.021)	(0.025)	(0.028)	(0.029)	(0.031)
100 to 150	0.018	-0.043*	0.518***	-0.046+	-0.004	0.025
	(0.013)	(0.018)	(0.022)	(0.024)	(0.026)	(0.027)
150 to 200	-0.025*	0.041**	-0.026	0.533***	-0.072**	-0.107***
	(0.012)	(0.016)	(0.019)	(0.021)	(0.022)	(0.024)
200 to 250	-0.002	-0.021	-0.019	-0.039*	0.480***	-0.100***
	(0.011)	(0.015)	(0.018)	(0.02)	(0.021)	(0.022)
250 to 300	-0.033**	-0.043**	-0.059***	-0.070***	-0.094***	0.330***
	(0.01)	(0.014)	(0.017)	(0.019)	(0.02)	(0.021)
Joint F test	92.3	122.3	107.9	112.4	96.8	44.4
Mean, Dep. Var.	0.044	0.095	0.145	0.193	0.227	0.256
B. Dependent Variable is Years Since First Containerization						
Depth is $\geq$ 30 feet, county is d_1 to d_2 km of any port, $t > 1956$						
0 to 50	25.591***	1.742	-1.373	-2.038	-6.203**	-0.5
	(1.141)	(1.577)	(1.888)	(2.12)	(2.223)	(2.341)
50 to 100	1.808*	31.261***	1.267	-1.688	2.356	-1.081
	(0.892)	(1.233)	(1.477)	(1.658)	(1.738)	(1.83)
100 to 150	1.082	-2.558*	31.083***	-2.757+	-0.238	1.5
	(0.786)	(1.086)	(1.3)	(1.46)	(1.531)	(1.612)
150 to 200	-1.476*	2.485**	-1.588	32.006***	-4.322**	-6.399***
	(0.691)	(0.954)	(1.143)	(1.283)	(1.345)	(1.416)
200 to 250	-0.148	-1.236	-1.113	-2.343*	28.817***	-6.028***
	(0.637)	(0.88)	(1.054)	(1.183)	(1.24)	(1.306)
250 to 300	-1.983**	-2.551**	-3.542***	-4.199***	-5.643***	19.802***
	(0.604)	(0.835)	(1)	(1.122)	(1.177)	(1.239)
Joint F test	92.3	122.3	107.9	112.4	96.8	44.4
Mean, Dep. Var.	2.7	5.7	8.7	11.6	13.6	15.4

Notes: Standard errors are below coefficients in parentheses. + indicates significance at the 10% level, * at the 5% level, ** at the 1% level, and *** at the 0.1 % level. All regression have 2,702 observations, and controls are as noted in Table 2b. Sources: See data appendix for details.

Table 4: Change in Log Population, 1950 to 2010, by Distance to Containerized Port

(a) Worldwide

	OLS Estimates				IV Estimates			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ever Cont. (0-100km)	-0.0613	0.0362	0.0817^{*}	0.1454^{***}	-0.8494^{***}	-0.5278^{***}	0.1695	0.2488^{**}
	(0.0656)	(0.0617)	(0.0491)	(0.0488)	(0.1473)	(0.1370)	(0.1078)	(0.1037)
Ever Cont. (100-200km)	-0.1910^{***}	-0.0252	0.0436	0.0183	-0.5693^{***}	-0.2839^{**}	0.2938^{**}	0.2343^{**}
	(0.0671)	(0.0637)	(0.0528)	(0.0511)	(0.1481)	(0.1414)	(0.1155)	(0.1084)
Ever Cont. (200-300km)	-0.1553^{**}	-0.0389	0.0929^{*}	0.0790	-1.0758^{***}	-0.6995^{***}	0.2958^{**}	0.2297^{*}
	(0.0661)	(0.0609)	(0.0564)	(0.0529)	(0.1558)	(0.1434)	(0.1355)	(0.1274)
Number of ports in 1953	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Country dummies	No	No	Yes	Yes	No	No	Yes	Yes
Log population in 1950	No	No	No	Yes	No	No	No	Yes

Sample: World cities over 50K in 1950.

Dependent variable: Change in log population between 1950 and 2010.

All specifications include a dummy variable for having a port within 300km in 1953.

The mean of the dependent variable is 1.54.

All regressions have 1051 observations.

Robust standard errors in parentheses. Stars denote significance levels: * 0.10 ** 0.05 *** 0.01.

Table 4: Change in Log Population, 1950 to 2010, by Distance to Containerized Port

(b) United States

	OLS Estimates				IV Estimates			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ County is d_1 to d_2 km from a containerized port								
0 to 50	0.311***	0.409***	0.241***	0.240***	0.347**	0.694***	0.349*	0.409*
	(0.068)	(0.078)	(0.069)	(0.07)	(0.127)	(0.192)	(0.171)	(0.177)
50 to 100	0.351***	0.295***	0.153**	0.152**	0.428***	0.354**	0.086	0.121
	(0.049)	(0.056)	(0.048)	(0.049)	(0.093)	(0.129)	(0.113)	(0.114)
100 to 150	0.138**	0.137**	0.101*	0.095*	0.154+	0.147	0.039	0.03
	(0.043)	(0.048)	(0.041)	(0.041)	(0.083)	(0.111)	(0.097)	(0.098)
150 to 200	0.081*	0.083*	0.078*	0.076*	0.094	0.122	0.163+	0.173*
	(0.039)	(0.042)	(0.036)	(0.036)	(0.08)	(0.098)	(0.085)	(0.088)
200 to 250	0.049	0.064	0.069+	0.057	0.150+	0.192+	0.287***	0.286**
	(0.038)	(0.041)	(0.035)	(0.035)	(0.081)	(0.098)	(0.086)	(0.088)
250 to 300	-0.007	0.007	0.026	0.016	0.039	0.125	0.211+	0.202
	(0.039)	(0.04)	(0.035)	(0.035)	(0.11)	(0.135)	(0.117)	(0.126)
Covariates								
Ever 1953 Port	x	x	x	x	x	x	x	x
Distance Bins to 1953 Port		x	x	x		x	x	x
Log Population, 1920-1950			x	x			x	x
1955 Int'l Trade & 1956 Manuf				x				x

Notes: Standard errors are below coefficients in parentheses. + indicates significance at the 10% level, * at the 5% level, ** at the 1% level, and *** at the 0.1 % level. The dependent variable mean is 0.376. See text for details on exact covariates. All regressions have 2,702 observations. Sources: See data appendix for details.

Table 5: Change in Log Population, 1950 to 2010, by Distance to Containerized Port and Years Since Containerization

(a) Worldwide

	OLS Estimates				IV Estimates			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Years since Cont. (0-100km)	-0.0071***	-0.0031*	0.0019	0.0040***	-0.0171***	-0.0131***	0.0047*	0.0067**
	(0.0017)	(0.0018)	(0.0015)	(0.0015)	(0.0029)	(0.0031)	(0.0028)	(0.0027)
Years since Cont. (100-200km)	-0.0098***	-0.0039**	0.0019	0.0012	-0.0129***	-0.0087**	0.0080***	0.0063**
	(0.0018)	(0.0019)	(0.0016)	(0.0015)	(0.0031)	(0.0034)	(0.0031)	(0.0029)
Years since Cont. (200-300km)	-0.0099***	-0.0053***	0.0017	0.0010	-0.0204***	-0.0162***	0.0067**	0.0050*
	(0.0017)	(0.0018)	(0.0016)	(0.0015)	(0.0031)	(0.0034)	(0.0032)	(0.0030)
Number of Ports in 1953	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Country dummies	No	No	Yes	Yes	No	No	Yes	Yes
Log Population in 1950	No	No	No	Yes	No	No	No	Yes

Sample: World cities over 50K in 1950.

Dependent variable: Change in log population between 1950 and 2010.

All specifications include a dummy variable for having a port within 300km in 1953.

The mean of the dependent variable is 1.54.

All regressions have 1051 observations.

Robust standard errors in parentheses. Stars denote significance levels: * 0.10 ** 0.05 *** 0.01.

Table 5: Change in Log Population, 1950 to 2010, by Distance to Containerized Port and Years Since Containerization

(b) United States

	OLS Estimates				IV Estimates			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ County is d_1 to d_2 km from a containerized port, Years Since First Cont.								
0 to 50	0.005***	0.007***	0.004***	0.004***	0.006**	0.012***	0.006*	0.007*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.003)	(0.003)	(0.003)
50 to 100	0.006***	0.005***	0.003**	0.003**	0.007***	0.006**	0.001	0.002
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
100 to 150	0.002**	0.002**	0.002*	0.002*	0.003+	0.002	0.001	0
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
150 to 200	0.001*	0.001*	0.001*	0.001*	0.002	0.002	0.003+	0.003*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
200 to 250	0.001	0.001	0.001+	0.001	0.003+	0.003+	0.005***	0.005**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
250 to 300	0	0	0	0	0.001	0.002	0.004+	0.003
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
Covariates								
Ever 1953 Port	x	x	x	x	x	x	x	x
Distance Bins to 1953 Port		x	x	x		x	x	x
Log Population, 1920-1950			x	x			x	x
1955 Int'l Trade & 1956 Manuf				x				x

Notes: Standard errors are below coefficients in parentheses. + indicates significance at the 10% level, * at the 5% level, ** at the 1% level, and *** at the 0.1 % level. The dependent variable mean is 0.376. See text for details on exact covariates. All regressions have 2,702 observations. Sources: See data appendix for details.

Table 6: Containerization Impact on Population Robust to Sample and Covariate Changes

(a) Worldwide

	Above 50K in 1950		Above 300K in 1950		Above 50K in 1950		Above 300K in 1950	
	(1) All	(2) Port cities	(3) All	(4) Port cities	(5) All	(6) Port cities	(7) All	(8) Port cities
Ever Cont. (0-100km)	0.2488**	0.2636**	0.2096	0.2012				
	(0.1037)	(0.1099)	(0.1665)	(0.1655)				
Ever Cont. (100-200km)	0.2343**	0.2192*	0.4099**	0.4623**				
	(0.1084)	(0.1186)	(0.1980)	(0.2224)				
Ever Cont. (200-300km)	0.2297*	0.2147	0.3712	0.3522				
	(0.1274)	(0.1415)	(0.2777)	(0.3075)				
Years since Cont. (0-100km)					0.0067**	0.0071**	0.0049	0.0051
					(0.0027)	(0.0028)	(0.0040)	(0.0041)
Years since Cont. (100-200km)					0.0063**	0.0057*	0.0112**	0.0118**
					(0.0029)	(0.0031)	(0.0050)	(0.0054)
Years since Cont. (200-300km)					0.0050*	0.0049	0.0058	0.0054
					(0.0030)	(0.0033)	(0.0054)	(0.0061)
Mean of Dep Variable	1.54	1.47	1.02	0.93	1.54	1.47	1.02	0.93
R-squared	0.68	0.71	0.78	0.79	0.68	0.72	0.79	0.80
Observations	1051	636	303	213	1051	636	303	213

Columns (1)-(2), (5)-(6): World cities over 50K in 1950.

Columns (3)-(4), (7)-(8): World cities over 300K in 1950.

Even numbered columns: Restrict sample to cities with a port within 300K in 1953.

Dependent variable: Change in log population between 1950 and 2010.

All specifications control for country fixed effects, log population in 1950, a dummy variable for having a port within 300km in 1953, and the number of ports within 300km in 1953.

Robust standard errors in parentheses. Stars denote significance levels: * 0.10 ** 0.05 *** 0.01.

Table 6: Containerization Impact on Population Robust to Sample and Covariate Changes

(b) United States

	Independent Variable is					
	Ever Containerized			Years Since First Containerization		
	Port Cities Only (1)	1950 Pop $\geq$ Median (2)	1950 Pop $<$ Median (3)	Port Cities Only (4)	1950 Pop $\geq$ Median (5)	1950 Pop $<$ Median (6)
Δ County is d_1 to d_2 km from a containerized port						
0 to 50	0.283+	0.193	0.908	0.005+	0.003	0.015
	(0.155)	(0.194)	(0.637)	(0.003)	(0.003)	(0.011)
50 to 100	0.032	0.16	0.156	0.001	0.003	0.003
	(0.099)	(0.145)	(0.233)	(0.002)	(0.002)	(0.004)
100 to 150	0	-0.012	0.003	0	0	0
	(0.083)	(0.119)	(0.204)	(0.001)	(0.002)	(0.003)
150 to 200	0.077	0.091	0.259	0.001	0.002	0.004
	(0.074)	(0.119)	(0.166)	(0.001)	(0.002)	(0.003)
200 to 250	0.224**	0.345***	0.122	0.004**	0.006***	0.002
	(0.074)	(0.103)	(0.176)	(0.001)	(0.002)	(0.003)
250 to 300	0.155	0.354*	0.066	0.003	0.006*	0.001
	(0.105)	(0.161)	(0.245)	(0.002)	(0.003)	(0.004)
Observations	1296	1351	1351	1296	1351	1351
Dep. Var. Mean	0.55	0.54	0.21	0.55	0.54	0.21

Notes: Standard errors are below coefficients in parentheses. + indicates significance at the 10% level, * at the 5% level, ** at the 1% level, and *** at the 0.1 % level. Controls are as noted in Table 2b. Sources: See data appendix for details.

A Data Appendix

A.1 Data Sources

We use data from a variety of sources. This appendix provides source information.

1. World Urbanization Prospects: The 2014 Revision
   These data include population counts for all urban agglomerations whose populations exceed 300,000 at any time between 1950 and 2010.
   - Downloaded from `http://esa.un.org/unpd/wup/CD-ROM/WUP2014_XLS_CD_FILES/WUP2014-F22-Cities_Over_300K_Annual.xls`
2. County Business Patterns
   These data include total employment, total number of establishments (with some variation in this definition over time), and total payroll.
   - 1956: Courtesy of Gilles Duranton and Matthew Turner. See Duranton et al. (2014) for source details. We collected a small number of additional counties that were missing from the Duranton and Turner data.
   - 1967 to 1985: U.S. National Archives, identifier 313576.
   - 1986 to 2011: U.S. Census Bureau. Downloaded from `https://www.census.gov/econ/cbp/download/`
3. Decennial Census: Population and employment data by county
   - 1950: ICPSR 02896, Historical, Demographic, Economic and Social Data: The United States, 1790-2002, Dataset 38: 1950 Census I (County and State)
   - 1960: ICPSR 02896, Historical, Demographic, Economic and Social Data: The United States, 1790-2002, Dataset 38: 1960 Census I (County and State)
   - 1970: ICPSR 8107, Census of Population and Housing, 1970: Summary Statistic File 4C – Population [Fourth Count]
   - 1980: ICPSR 8071, Census of Population and Housing, 1980: Summary Tape File 3A
   - 1990: ICPSR 9782, Census of Population and Housing, 1990: Summary Tape File 3A
   - 2000: ICPSR 13342, Census of Population and Housing, 2000: Summary File 3
   - 2010: U.S. Census Bureau, 2010 Decennial Census Summary File 1, Downloaded from `http://www2.census.gov/census_2010/04-Summary_File_1/`
4. Port Universe and Depth
   - We use these documents to establish the population of ports in any given year.

- 1953: National Geospatial Intelligence Agency (1953)
- 2015: National Geospatial Intelligence Agency (2015)

5. Port Containerization Adoption Year
    - 1956–2010: *Containerisation International Yearbook* for 1968 and 1970–2010
6. Port Volume: Total imports and exports by port
    - 1948: United States Foreign Trade, January-December 1949: Water-borne Trade by United States Port, 1949, Washington, D.C.: U.S. Department of Commerce, Bureau of the Census. FT 972.
    - 1955: United States Waterborne Foreign Trade, 1955, Washington, D.C. : U.S. Dept. of Commerce, Bureau of the Census. FT 985.
    - 2008: Containerisation International yearbook 2010, pp. 8–11.

A.2 Data Choices

1. U.S. County Sample

    We drop XX counties where land area changes are greater than 35 percent. These are: [list here].

    In a future draft, we will provide a list of the county groupings we use to make the 1950 and 2010 counties geographically compatible.

    Alaska and Hawaii were not states in 1950. We omit Alaska from our sample, because in 1950 it has only judicial districts, which do not correspond to modern counties. We keep Hawaii, where the 1950 borders are relatively equivalent to modern counties. We also keep Washington, DC, in all years.

2. Ports

    We determine the universe of ports from the 1953 *World Port Index*.

Appendix Figure 1: Sample Construction by Distance to Port

(a) Counties within 50 kilometers of port

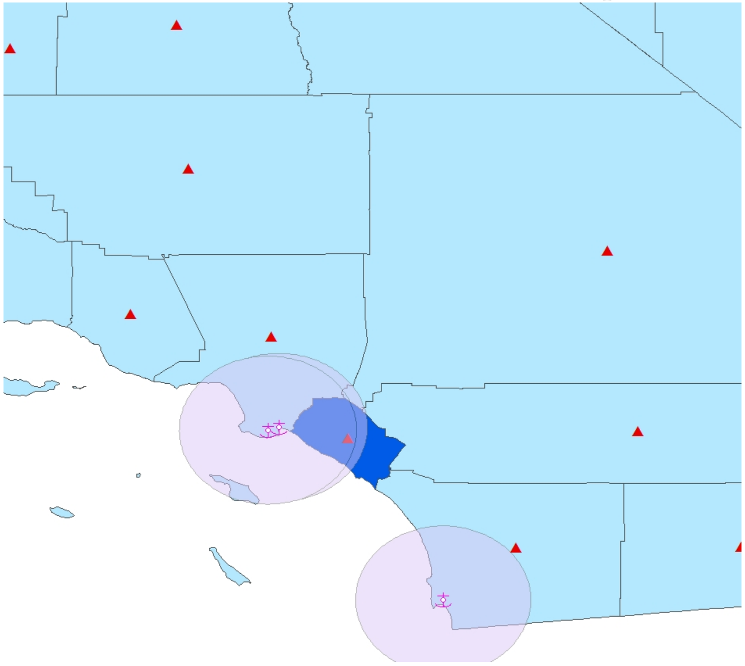

(b) Counties within 50 to 100 kilometers of port

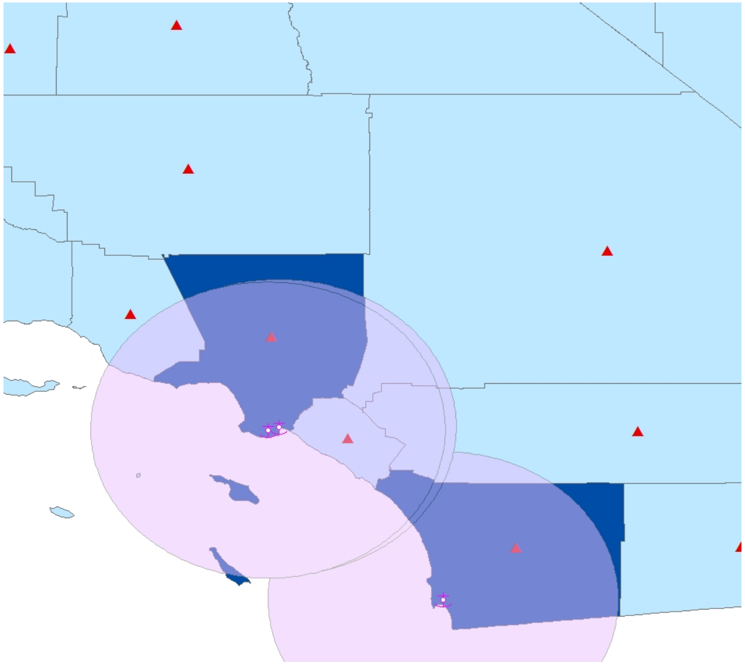

Notes: Blue polygons are counties, and red triangles are the geographic county centers (centroids). The pink anchors are ports, and the grey circles show rings of 50 (top figure) and 100 (bottom figure) from the ports.
Source: See data appendix.

Appendix Table 1: Pre-Containerization International Trade by Distance to Port, United States

	Distance to Containerized Port							
	0 to 50	50 to 100	100 to 150	150 to 200	200 to 250	250 to 300	Ever Cont.	Never Cont.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1955: $ Millions of Int'l Trade in All Ports by row's distance bin								
0 to 50	960.94	136.41	266.73	159.07	72.69	107.75	103.49	0.02
	[2084.54]	[713.37]	[1229.58]	[1011.12]	[634.61]	[831.15]	[747.27]	[0.53]
50 to 100	263.24	683.12	249.84	172.95	197.05	124.46	154.96	0.04
	[1136.6]	[1716.36]	[1120.74]	[818.1]	[1086.29]	[746.85]	[863.98]	[0.61]
100 to 150	710.68	334.99	633.56	278.33	194.41	140.75	224.27	0.19
	[1927.38]	[1216.6]	[1564.15]	[1057.4]	[830.96]	[697.41]	[985.61]	[3.33]
150 to 200	332.37	371.02	337.96	586.49	267.46	252.47	267.25	0.12
	[1221.24]	[1143.06]	[1213.83]	[1483.09]	[1059.51]	[1040.41]	[1037.56]	[1.12]
200 to 250	261.98	360.8	348.85	308.51	554.01	283.61	314.91	0.55
	[1018.96]	[1297.78]	[1196.64]	[1153.69]	[1466.17]	[1121.64]	[1161.58]	[7.39]
250 to 300	549.15	389.99	471.95	385.28	361.86	610.49	389.3	1.12
	[1794.07]	[1382.2]	[1578.65]	[1289.51]	[1285.22]	[1553.58]	[1302.81]	[10.28]
1948: $ Millions of Int'l Trade in All Ports by row's distance bin								
0 to 50	812.24	109.95	228.76	138.98	60.69	93.11	87.33	0.01
	[1914.6]	[641.93]	[1118.6]	[922.04]	[575.91]	[756.12]	[678.29]	[0.34]
50 to 100	215.96	554.39	201.69	137.78	163.56	99.9	125.45	0.02
	[1005.68]	[1548.83]	[999.88]	[730.29]	[976.03]	[671.34]	[772.43]	[0.4]
100 to 150	591.27	271.95	502.97	223.56	149.21	105.63	177.95	0.12
	[1719.97]	[1093.62]	[1402.64]	[950.31]	[735.69]	[619.91]	[877.08]	[2.66]
150 to 200	261.62	286.24	272.87	462.03	215.58	201.07	209.54	0.05
	[1111.75]	[1016.17]	[1090.86]	[1334.58]	[950.47]	[930.03]	[926.23]	[0.64]
200 to 250	198.16	283.27	274.84	249.66	437.89	228.65	250.17	0.31
	[917.23]	[1162.1]	[1070.82]	[1032.86]	[1319.82]	[1005.19]	[1040.76]	[3.93]
250 to 300	452.27	317.42	387.24	309.64	293.65	485.82	311.54	0.66
	[1593.87]	[1244.47]	[1418.54]	[1157.74]	[1143.53]	[1398.74]	[1167.75]	[5.12]

Appendix Table 2: Containerization More Likely in Counties Near Deeper Ports: Alternative Depth Measures

	Dependent Variable is							
	Ever Containerized				Years Since First Containerization			
	Anchorage Depth		Channel Depth		Anchorage Depth		Channel Depth	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Depth in feet is								
30-35	0.868***		0.909***		35.869***		38.308***	
	(0.047)		(0.052)		(2.376)		(2.653)	
30-35	0.958***		0.975***		43.499***		39.255***	
	(0.049)		(0.056)		(2.458)		(2.843)	
25-30	0.961***		0.951***		40.453***		40.849***	
	(0.046)		(0.051)		(2.332)		(2.591)	
20-25	0.892***		0.897***		37.382***		36.253***	
	(0.049)		(0.051)		(2.447)		(2.613)	
15-20	0.396***		0.913***		15.494***		36.267***	
	(0.055)		(0.052)		(2.756)		(2.647)	
10-15	0.930***		-0.005		30.493***		0.085	
	(0.061)		(0.1)		(3.075)		(5.11)	
	.		-0.004		.		-0.411	
	.		(0.071)		.		(3.617)	
1{Depth $\geq$ 30 Feet}		0.211***		0.082***		10.324***		5.451***
		(0.016)		(0.011)		(0.802)		(0.523)
Mean, Dependent Variable	0.447	0.447	0.447	0.447	19.85	19.85	19.85	19.85
R-squared	0.909	0.889	0.909	0.885	0.888	0.87	0.887	0.867
Observations	2702	2702	2702	2702	2702	2702	2702	2702
F for Excluded Instrument(s)	130	170	115	58	106	166	84	109
Increase in R^2 due to instrument	0.026	0.007	0.027	0.003	0.027	0.008	0.025	0.005

Notes: We report standard errors below coefficients in parentheses. + indicates significance at the 10% level, * at the 5% level, ** at the 1% level, and *** at the 0.1 % level. Port depth is the depth of the deepest port within 300 km. All specifications control for a dummy for ever being within 300 km of a 1953 port, the number of 1953 ports in each of six distance bins to 300 km, log population 1920 to 1950, 1956 manufacturing share, and the total value of waterborne international trade in each of six distance bins to 300 km. Results using wharf depth are in Table 2b.

Source: See data appendix.

Appendix Table 3: Only Cities Within 300 km of Ports: Instrument Remains Strong

(a) Worldwide

	Ever Cont.			Years since Cont.		
	(1) 0 to 100	(2) 100 to 200	(3) 200 to 300	(4) 0 to 100	(5) 100 to 200	(6) 200 to 300
30+ ft. Depth (0-100km)	0.543***	-0.094**	-0.045	20.690***	-3.591**	-2.898**
	(0.039)	(0.039)	(0.041)	(1.506)	(1.393)	(1.457)
30+ ft. Depth (100-200km)	-0.087**	0.486***	-0.083**	-2.755**	17.839***	-3.238**
	(0.038)	(0.039)	(0.041)	(1.367)	(1.484)	(1.420)
30+ ft. Depth (200-300km)	-0.093**	-0.100**	0.459***	-2.809*	-3.040*	18.401***
	(0.043)	(0.044)	(0.044)	(1.561)	(1.556)	(1.637)
Mean, Dependent Variable	0.57	0.57	0.59	19.00	18.07	19.48
R-squared	0.50	0.50	0.51	0.56	0.55	0.56
F Stat Excluded Instrument(s)	83.34	63.25	50.01	80.11	59.52	68.08

Sample: World cities over 50K in 1950, excluding cities without a port within 300km in 1953.
Dependent variable (1)-(3): Adoption of containerization between d1 and d2 km.
Dependent variable (4)-(6): Number of years since the adoption of containerization between d1 and d2 km.
Port depth is the depth of the deepest port within 300km.
All specifications control for country fixed effects, log population in 1950,
a dummy variable for having a port within 300km in 1953, and the number
of ports within 300km in 1953.
All regressions have 1051 observations.
Robust standard errors in parentheses. Stars denote significance levels: * 0.10 ** 0.05 *** 0.01.

Appendix Table 3: Only Cities Within 300 km of Ports: Instrument Remains Strong

(b) United States

	Dependent Variable is			
	Ever Containerized		Years Since First Cont.	
	(1)	(2)	(3)	(4)
Port Depth in feet is				
40 and over	1.038***		46.563***	
	(0.092)		(4.521)	
35-40	1.004***		45.742***	
	(0.091)		(4.459)	
30-35	0.890***		35.789***	
	(0.091)		(4.448)	
25-30	0.956***		41.149***	
	(0.091)		(4.451)	
20-25	0.704***		27.742***	
	(0.095)		(4.653)	
15-20	-0.017		0.647	
	(0.155)		(7.617)	
10-15	0.01		0.091	
	(0.122)		(5.973)	
1{Depth $\geq$ 30 Feet}		0.079***		3.830***
		(0.016)		(0.78)
Mean, Dependent Variable	0.933	0.933	41.384	41.384
R-squared	0.256	0.055	0.366	0.19
Observations	1296	1296	1296	1296
F for Excluded Instrument(s)	54	25	55	24
Increase in R^2 due to instrument	0.219	0.018	0.191	0.015

Notes: We report standard errors below coefficients in parentheses. + indicates significance at the 10% level, * at the 5% level, ** at the 1% level, and *** at the 0.1 % level. Port depth is the depth of the deepest port within 300 km. All specifications control for a dummy for ever being within 300 km of a 1953 port, the number of 1953 ports in each of six distance bins to 300 km, log population 1920 to 1950, 1956 manufacturing share, and the total value of waterborne international trade in each of six distance bins to 300 km.
Source: See data appendix.

Appendix Table 4: Specific Depth Cut-off For Instrument Not Binding

	Dependent Variable: County is d_1 to d_2 km of containerized port					
	0 to 50	50 to 100	100 to 150	150 to 200	200 to 250	250 to 300
	(1)	(2)	(3)	(4)	(5)	(6)
Depth is $\geq$ 25 feet, county is d_1 to d_2 km of any port, $t > 1956$						
0 to 50	0.478***	0.221***	-0.400***	-0.109***	-0.025	-0.014
	(0.016)	(0.022)	(0.025)	(0.027)	(0.028)	(0.029)
50 to 100	0.057***	0.512***	0.139***	-0.415***	-0.071**	-0.051+
	(0.015)	(0.02)	(0.022)	(0.025)	(0.025)	(0.027)
100 to 150	0.024+	0.100***	0.545***	0.125***	-0.532***	-0.096***
	(0.014)	(0.02)	(0.022)	(0.024)	(0.025)	(0.026)
150 to 200	0.022	0.074***	0.107***	0.549***	0.115***	-0.481***
	(0.015)	(0.021)	(0.023)	(0.026)	(0.026)	(0.028)
200 to 250	-0.007	0.007	0.060***	0.117***	0.670***	0.233***
	(0.011)	(0.016)	(0.018)	(0.019)	(0.02)	(0.021)
250 to 300	-0.01	0	0.003	0.086***	0.098***	0.660***
	(0.012)	(0.017)	(0.019)	(0.021)	(0.021)	(0.022)
Joint F test	216.9	222.8	221.2	203.1	197.8	153.3
Depth is $\geq$ 35 feet, county is d_1 to d_2 km of any port, $t > 1956$						
0 to 50	0.624***	0.168***	-0.600***	-0.119*	0.024	-0.105*
	(0.027)	(0.038)	(0.044)	(0.047)	(0.05)	(0.051)
50 to 100	0.130***	0.657***	0.055	-0.645***	-0.153***	-0.026
	(0.025)	(0.035)	(0.04)	(0.044)	(0.046)	(0.047)
100 to 150	0.007	0.174***	0.532***	-0.127**	-0.531***	-0.104*
	(0.025)	(0.035)	(0.04)	(0.043)	(0.045)	(0.046)
150 to 200	0.01	0.026	0.126**	0.617***	-0.107*	-0.492***
	(0.024)	(0.034)	(0.04)	(0.043)	(0.045)	(0.046)
200 to 250	0.009	-0.015	0.121***	0.245***	0.500***	0.121***
	(0.018)	(0.026)	(0.03)	(0.032)	(0.034)	(0.035)
250 to 300	0.01	0.068**	-0.006	0.045	0.224***	0.520***
	(0.018)	(0.025)	(0.029)	(0.031)	(0.033)	(0.034)
Joint F test	127.9	126.9	84.4	80.4	44.5	42.4
Depth is $\geq$ 40 feet, county is d_1 to d_2 km of any port, $t > 1956$						
0 to 50	0.683***	-0.075	-0.574***	-0.047	-0.05	-0.194*
	(0.048)	(0.068)	(0.077)	(0.083)	(0.085)	(0.087)
50 to 100	-0.002	0.591***	0.033	-0.393***	-0.089	0.04
	(0.04)	(0.057)	(0.064)	(0.07)	(0.071)	(0.072)
100 to 150	-0.085*	0.169**	0.547***	-0.198**	-0.317***	0.01
	(0.038)	(0.054)	(0.061)	(0.066)	(0.068)	(0.069)
150 to 200	0.090*	0.079	0.024	0.349***	-0.042	-0.352***
	(0.035)	(0.05)	(0.056)	(0.061)	(0.062)	(0.063)
200 to 250	0.052+	-0.049	-0.068	0.231***	0.376***	0.076
	(0.027)	(0.039)	(0.044)	(0.047)	(0.049)	(0.049)
250 to 300	0.012	0.008	-0.016	0.089*	0.135**	0.422***
	(0.025)	(0.035)	(0.04)	(0.043)	(0.044)	(0.045)
Joint F test	50.98	42.024	19.132	17.604	11.014	16.29

Notes: We report standard errors below coefficients in parentheses. + indicates significance at the 10% level, * at the 5% level, ** at the 1% level, and *** at the 0.1 % level. Port depth is the depth of the deepest port within 300 km. Results using wharf depth are in Table 3b. Source: See data appendix.