Economic Activity Across Space: A Supply and Demand Approach*

Treb Allen and Costas Arkolakis

March 2023

Abstract

What do recent advances in economic geography teach us about the spatial distribution of economic activity? We show that the equilibrium distribution of economic activity can be determined simply by the intersection of labor supply and demand curves. We discuss how to estimate these curves and highlight the importance of global geography – i.e. the connections between locations through the trading network – in determining how various policy relevant changes to geography shape the spatial economy.

1 Introduction

The spatial distribution of people is incredibly concentrated: 8 percent of the US population lives in the ten largest US cities, and but those cities take up less than 0.1 percent of total US land area. Why this concentration? More generally, what determines the distribution of people and economic activity across space? And how can economic policies affect the spatial distribution of economic activity? This essay will show that these questions can be answered through the familiar lens of supply and demand curves.

We begin by applying this intuition to the well-known Rosen-Roback (Rosen, 1979; Roback, 1982) framework. But as we will discuss, the distribution of economic activity

*Treb Allen is Distinguished Associate Professor of Economics and Globalization, Dartmouth College, Hanover, New Hampshire. Costas Arkolakis is Professor of Economics, Yale University, New Haven, Connecticut. Their email addresses are treb.allen@dartmouth.edu and costas.arkolakis@yale.edu. We thank our editors Erik Hurst, Nina Pavcnik, Timothy Taylor, and Heidi Williams as well as Steve Redding, Anson Soderbery, and Sayantan Roy for exceptionally helpful feedback and Saptarshi Majumdar for excellent research assistance. The authors acknowledge support by the National Science Foundation under grants SES-1658838 and SES-1658875.

in this early spatial model depends only on local geography, not on what happens to other regions. For example, a change in one location – say, a large infrastructure investment that improves its productivity – is predicted to have an identical impact on all other locations, regardless of where they are. Thus, intuitive spatial features like where a location is located on a map and who its neighbors are entirely absent: it is a spatial model where space does not matter.

In reality, spatial linkages create rich interactions between locations. One implication of these interactions is that a large infrastructure investment that improves the productivity in one location will have greater impacts on close by locations than locations further away. To account for such spatial linkages, we extend the intuition of the Rosen-Roback model to modern economic geography frameworks where locations are connected through the flow of goods, based on our earlier work in Allen and Arkolakis (2014). In this framework, the economic fate of a location depends not only on its own "local" geography but also on the local geography of its neighbors, the effect of which is mediated by the strength of the economic ties, creating a "global geography". Despite this added complexity, we show the same tools based on supply and demand used to understand predictions of the earlier Rosen-Roback framework extend readily to a globally integrated world.

This globally integrated framework can be applied to understand both the direct and indirect impacts of real world economic policies that change either the local or global geography. We discuss how the framework can be applied to spatial data, while also highlighting the most common pitfalls and offering strategies for traversing them. Finally, we provide a brief overview of the many ways in which this framework has been applied thus to understanding the spatial distribution of economic activity, as well as pointing out several interesting and as-yet unexplored questions for future researchers. To keep the discussion as straightforward and accessible as possible, we relegate all mathematical details and derivations to the Appendix. We also provide a companion Matlab toolkit to help researchers apply these techniques on their own.

2 Understanding the spatial distribution of economic activity through the lens of supply and demand

We now discuss the Rosen-Roback framework. Consider a world comprising many different locations. These locations each have their own "local" geography. The "local" geography of a location includes a whole host of things, from natural geographic features like the climate, elevation, natural beauty, etc. to other less tangible characteristics of a location like the quality of its political institutions. Local geography can affect the spatial distribution of economic activity in two ways. First, it can affect the desire of people to live in a location and hence labor supply; we will call such factors "amenities." Second, it can affect how productive people are in a location and hence labor demand; we will call such factors "productivities."

Figure 1 illustrates the spatial equilibrium that results for labor demand and labor supply in this market. The labor market envisaged in this Rosen-Roback approach is one defined by location, rather than by the specific skills or sectors of workers: we think about the supply and demand for all workers in Detroit rather than the supply and demand for nurses or auto mechanics.

Let us first examine the labor demand curve more closely. Wherever people choose to live, they earn a wage from producing a good and then use that wage to buy goods and services. Let us assume that the wage they earn in any location i depends on two things: (1) the number of people living in that location; and (2) the productivities of that location. The result is a labor demand curve:

$$\ln w_i = \varepsilon^D \ln L_i + \ln C_i^D. \tag{1}$$

In this relationship, the terms for wages and quantity of labor are expressed in log terms, and so ε^D is the demand elasticity. C_i^D is the local productivity in region i that arises from its local geography. The local productivity may capture, for example, how productive the factors are in location i or the relative cost of capital in a location.

The elasticity of demand is typically assumed to be negative, such that the labor demand function in Figure 1 is downward-sloping. The economic intuition behind this slope is often based on assuming decreasing returns to scale in production of the good or simply the presence of a fixed factor such as capital (for example, see Kline and Moretti (2014) and Donaldson and Hornbeck (2016)). In other words, there is diminishing marginal product for each additional unit of labor added in the location. Thus, as the population of a location increases, each additional worker is less and less productive, causing the wage to fall. But other scenarios are possible. For example, the presence of external economies also can affect the slope of the demand function. If more workers in a location result in everyone being more productive, the labor demand curve can become more elastic; if these external economies are sufficiently strong, the demand curve may even slope upwards. This situation can lead to outcomes like multiple equilibria or like "black hole" equilibria where everyone lives in one location (for discussion, see Krugman (1991) and Fujita, Krugman, and Venables (1999)). While such scenarios have academic interest, in what follows we will stick with the more common (and, arguably, more empirically relevant) case of a downward-sloping demand curve.

If people each choose their place of residence to be as happy as possible, what makes people happy in this framework? Two things: higher consumption (so, all else equal, workers prefer higher real wages) and living somewhere nice (that is, a place with high amenities). In a model where everyone is identical, all inhabited locations must make people equally as happy. If prices are the same everywhere (so that the real wage is the nominal wage) and the amenity value of a location depends in part on how many other people live there, then workers' indifference across all inhabited locations generates this labor supply curve:

$$\ln w_i = \varepsilon^S \ln L_i - \ln C_i^S. \tag{2}$$

Again, the left-hand side of the equation is the wage for each worker in the region i and L_i is the number of workers in the region. Because wages and the quantity of workers are expressed in logs, where ε^S is the elasticity of labor supply. C_i^S is the local amenity in region i, e.g. better parks or planetaria.[1]

Economists usually think of a supply curve as sloping upward, as the labor supply curve is shown in Figure 1. A common underlying assumption in this setting is that the supply curve will slope up as long as more people in a location make each individual less happy; for example, the presence of a housing market where a higher population drives up housing prices and rents, or the existence of idiosyncratic preferences where a higher population means the marginal resident's match quality is worse can also lead to upward sloping labor supply curves.[2] It is theoretically possible for the labor supply curve to slope downward (and issues of multiplicity and black holes to arise) if the amenity value of a location is increasing in its population, perhaps because of greater investments in public goods or greater variety in consumables in that specific location, but we will set that possibility aside here.

In this model, the equilibrium of economic activity – that is, the population and wage in a specific location – arises from combining the labor demand and labor supply curves. The spatial equilibrium is highlighted at point A in Figure 1.

To see how the "local" geography shapes the spatial equilibrium, consider a simple counterfactual scenario where the amenity value of residing in a location improved. For example, suppose the advent of air conditioning technology made the hot climate of the U.S. Southwest less oppressive. An improvement in amenities shifts outward the labor supply curve, moving the equilibrium from point A to point B in Figure 1. The population in the location

[1] See the online Appendix A.1 for a particular micro-foundation that delivers the specific labor demand and labor supply functions shown here.

[2] For a discussion of heterogeneous preferences and housing market see Helpman (1998); Allen and Arkolakis (2014); Redding (2016); Ahlfeldt, Redding, Sturm, and Wolf (2015).

increases, but its wage declines. The U.S. Southwest is now a better place to live, but the influx of workers depresses the wages.

The fact that we can analyze each location separately, depending on the amenity shock it receives, illustrates the somewhat paradoxical nature of the Rosen-Roback framework. It is a spatial model, but the distribution of economic activity depends only on local geography, not on what happens to other regions. Intuitive spatial features like where a location is located on a map and who its neighbors are entirely absent: it is a spatial model where space does not matter. By looking at one location at a time, it does not consider economic linkages between those locations. [3] Taking such linkages into account will create the concept of "global" geography which we introduce and analyze next.

3 The role of global geography in the spatial distribution of economic activity

Different locations can be linked with each other in many ways: people may live in one location and work in another; people may migrate from one location; people may talk with each other, leading to the spatial diffusion of ideas; and so on. But perhaps the most obvious spatial linkage occurs through the flow of goods. Much of what an individual consumes is produced in another location: according to the 2017 United States Commodity Flow Survey (CFS, 2017),[4] most freight shipments crossed state boundaries, with only 22% of the value of freight destined for a state also originating in the same state. Moreover, the pattern of trade flows are far from uniform. As panel (a) of Figure 2 highlights using the same data, nearby states trade more with each other while the total volume of trade increases with the size of the trading partners, a phenomenon originally observed in international trade flows and oftentimes referred to as "gravity" (Anderson, 2010; Head and Mayer, 2013).

How does incorporating such spatial linkages affect the spatial equilibrium? It turns out that much of the basic intuition above remains: in particular, we can still analyze the spatial equilibrium using the familiar techniques of supply and demand, albeit now augmented with a concept of both "local" and "global" geographies.

[3] In the Rosen-Roback framework, a change in the local geography in one location can have aggregate general equilibrium effects on, say, the price of capital. But such general equilibrium effects affect all locations equally and hence do not affect the spatial distribution of economic activity.

[4] The Commodity Flow Survey is conducted by the U.S. government and is the primary source of data on within-U.S. trade flows. In general, it is difficult to measure intra-country trade flows, making analysis of within-country trade difficult, although notable exceptions include work in Canada (Anderson and Van Wincoop, 2003; McCallum, 1995), India (Donaldson, 2018), and the Philippines (Allen, 2014).

The global geography

The model discussed below is based on prior work, Allen and Arkolakis (2014), but variations of this spatial framework, with equivalent or similar mathematical formulations have recently been used in a variety of frameworks.[5] The setup retains the same features as above but now we introduce a key distinction: goods are no longer costlessly traded. There are trade relationships between different locations, governed by the presence of spatial frictions.

These spatial frictions can be described as the *economic distance* between regions i and j. Conceptually, economic distance is proportional to the value of trade flows between two locations (conditional on origin and destination fixed effects). There are many possible factors that influence the economic distance between locations – whether they speak the same language, share the same legal systems, share similar cultural heritages, etc. But one of the most important contributors to economic distance is simply the geographic distance between any two locations. Indeed, one of the most robust empirical relationships in all of economics is that trade flows between locations are roughly inversely proportional to the geographic distance between them (for discussion, see Disdier and Head (2008) and Chaney (2018)). Put another way, a very good start to measuring "economic distance" is simply with geographic distance.

When spatial frictions exist and goods are no longer costlessly traded, two things change. First, the prices of the goods produced by workers in a location depends in part on how nearby the consumers of those products are. The closer the consumers are, the more demand for their products, and the higher the price (and hence the higher the wage) that the workers can obtain. This *outward market access* affects the labor demand curve of a location. Second, the price of goods purchased by consumers in a location depend in part on how nearby the producers of those products are. The closer the producers, the lower the price for those products, and the higher the real wage of the consumers. This *inward market access* acts as a shifter to the labor supply curve of a location. Together, the outward and inward market accesses comprise the *global geography* of a location.

Together, the outward and inward market accesses comprise the *global geography* of a location. Following Anderson and Van Wincoop (2003); Redding and Venables (2004), the outward market access (MA_i^{out}) can mathematically be expressed as:

$$MA_i^{out} = \sum_j T_{ij} \times \frac{Y_j}{MA_j^{in}}, \tag{3}$$

[5]See for example Redding (2016); Donaldson and Hornbeck (2016); Allen, Arkolakis, and Takahashi (2020); Faber and Gaubert (2019); Eckert and Peters (2022). Redding and Rossi-Hansberg (2017) offer a comprehensive review of the quantitative spatial framework.

where T_{ij} is the inverse of economic distance between two locations and $Y_j = w_j L_j$ is the total income of location j. Intuitively, outward market access summarizes the selling potential of a market, summarizing how well a region is connected to other locations. For example, New Jersey has a high outward market access because there are lots of potential consumers of its products in its neighboring states of New York and Pennsylvania. It is greater when its neighboring locations are closer (i.e. inverse economic distance T_{ij} is greater), which is especially beneficial when those neighboring locations are richer (i.e. have higher Y_j) or have worse alternatives for buying their own goods (i.e. have lower MA_j^{in}).

Inward market access, in turn, is similarly defined as the capacity of locations to buy from other locations:

$$MA_j^{in} = \sum_i T_{ij} \times \frac{Y_i}{MA_i^{out}}. \tag{4}$$

For example, New Jersey also has high inward market access because it is able to purchase its goods from nearby large producers. Like outward market access, inward market access is greater the smaller the economic distance to other locations, and again this matters more when nearby locations either produce a lot (i.e. have higher Y_i) or have poor alternatives for selling their goods (i.e. have a lower MA_i^{out}). Outward and inward market accesses are obviously quite closely related and, indeed, will be the proportional to each other in the special case when economic distances are the same in both directions. Note, however, that the economic distance that matters for inward market access is the one in which a location is the destination, whereas for outward market access, the economic distance that matters is the one in which the location is the origin. As a result, when economic distances are not the same in both directions, the inward and outward market accesses will generally be different.

The global geography summarizes how each location depends on economic activity in all other locations, where closer locations are given greater weights. These algebraic formulations highlight that inward and outward market accesses are intertwined, with each dependent in part on the other. Despite this interdependence, it is straightforward to solve for both the market access measures as long one observes the income in each location and the economic distances between locations. The companion Matlab code available as an appendix to this paper provides a convenient algorithm for doing so.

Panel (b) of Figure 2 depicts the (outward) market access for each U.S. states, where we proxy the inverse economic distance T_{ij} with inverse of geographic distance, measured as the distance (as the crow flies) between the geographic center of each state. States with high economic output that are close to other states with high output such as those in the Northeast have good market access; states with less economic output that are far away from states with higher economic output such as Montana have poor market access. As we will discuss in the next main section, an appealing feature of this framework is that the inverse economic distance can also be measured more explicitly with a combination of observed bilateral trade flows and observed bilateral geographic characteristics such as distance or time of travel.

The global spatial equilibrium

It turns out the global spatial equilibrium with spatial linkages can be analyzed using labor supply and demand curves just as in the local spatial equilibrium above. Now, however, supply and demand will not only depend on local geography, but also on global geography. In particular, the labor demand now also depends on outward market access MA_i^{out}, becoming:

$$\ln w_i = \varepsilon_{local}^D \ln L_i + \varepsilon_{global}^D \ln MA_i^{out} + \ln C_i^D. \tag{5}$$

Better outward market access acts analogously to better local productivities, C_i^D, shifting the demand curve for local labor outwards with an elasticity ε_{global}^D. That elasticity is greater the less substitutable the goods produced in i are with goods produced elsewhere in the world.

Similarly, labor supply now depends on inward market access MA_i^{in}, becoming:

$$\ln w_i = \varepsilon_{local}^S \ln L_i + \varepsilon_{global}^S \ln MA_i^{in} - \ln C_i^S. \tag{6}$$

Better inward market access acts analogously to better local amenities C_i^S, shifting the supply curve for labor outwards with an elasticity ε_{global}^S, which again is larger the less substitutable goods produced in different locations are with each other.

The two limiting cases deserve special mention. When $\varepsilon_{local}^S \to \infty$, the local population is invariant to changes in economic conditions, whereas when $\varepsilon_{local}^S \to 0$ the labor supply is infinitely elastic to local economic conditions. These special cases correspond to important cases in the literature, as we will discuss below.

Given the global geography, the global spatial equilibrium is determined just as in the local spatial equilibrium above: find the wage and population in each location that equates supply with demand; point A on panel (a) of Figure 3 depicts such an equilibrium.

So what has changed in the global spatial equilibrium? The crucial insight is that *the global geography in one location depends on the spatial equilibria in all other locations*. If something changes about the local geography anywhere in the world, it will affect the global geography everywhere in the world, although it will affect nearby locations more than locations far away. Hence, the global geography puts space back into the spatial economy.

To illustrate this global spatial equilibrium, let us return to the example above. Suppose that air conditioning is invented, which makes some hot and previously inhospitable location i much more hospitable, raising the amenity of living there. Again, this innovation will shift outward labor supply curve in location i to point B in panel a of Figure 3, increasing the population in location i and reducing the wages. The story does not end here, as this change in population and wages will affect the global geography. As long the as the elasticity of local demand is greater than -1, the income Y_i of location i will increase, raising both the inward and outward market access and resulting in an additional shift outward to both the labor demand and labor supply curves. This additional global effect further increases the population in location i and mitigates the downward fall in wages, as illustrated in point C in panel (a) of Figure 3.

At the same time, changes in the economic activity in location i affect the global geography of other locations. Consider a neighboring location j initially in equilibrium, as illustrated by point A in panel (b) of Figure 3. Because the income of location its neighbor i has improved, both its supply and demand curves will shift outwards as well. Intuitively, the greater nearby economic activity both increases the demand for the goods produced in j and increases the supply of goods consumed in j. As a result, the population in j increases too (and its wages rise), changing its equilibrium to point C in panel (b) of Figure 3, despite there being no change in its own local geography.[6]

But won't changes in the economic activity in location j have subsequent impacts on the global geography in all other locations? And won't those changes have even further impacts on the global geography, *ad infinitum*? Yes and yes: indeed, this infinite feedback loop between the global geography in every location is part of what makes the global spatial equilibrium so interesting to study. In reality, point C in panels (a) and (b) of Figure 3 represents the limit of the infinite sequence of these adjustments of each locations' global geography to adjustments made in the global geography everywhere else. Indeed, this iterative process is what both the algorithm for calculating the equilibrium change in market accesses in the companion Matlab code and many tools for studying the mathematical properties of the equilibrium system is based upon.[7]

[6]Whether nominal wages rise or fall – that is, whether outward or inward market access increases more – depends on the choice of the numeraire. Here we set mean wages equal to one as the numeraire, so falling wages in location i must be offset by rising wages elsewhere.

[7]In the special case where the augmented labor supply curve is infinitely elastic, the local and global demand elasticities are equal in magnitude, and the inverse economic distances are symmetric, the equilibrium global economy is one in which the wages and populations of each location are (log) proportional to the *eigenvector centrality* of a location in the network defined by the world geography (that is, by the combination of the economic distances, productivities, and amenities). Higher eigenvector centrality means that a node

Having shown how one can determine the global spatial equilibrium through the use of supply and demand curves, we now turn to describing the process through which this framework can be combined with spatial data to assess the impact of changes in geography on the real world spatial distribution of economic activity.

4 Estimating labor supply and demand

In the previous section, we saw how a supply and demand framework can be used to understand how changes in the geography affects the distribution of economic activity across spatially connected locations. One of the most attractive aspects of the global spatial framework described above is its ability to seamlessly integrate with readily available spatial data. In this section, we describe this interplay between theory and data.

Spatial Economic Data: Local and Linkages

We focus here on two types of spatial data: data on the local economic activity of a location; and data on the strength of economics linkages between locations across space.

Suppose that a researcher can observe in the data how many people reside in a certain location L_i and the total income of a location Y_i. Indeed, such data are readily available; for example, in the United States, population data and income data at the county level can be constructed from the decennial Census going back to the year 1840. The IPUMS National Historical Geographic Information Systems (Manson, 2020) has provided an enormous public good in assembling these data and making them publicly available. Even in parts of the globe where spatially disaggregated income data is not readily available, one can proxy for economic activity using satellite data on the intensity of lights at nighttime, a practice pioneered by Henderson, Storeygard, and Weil (2012) and summarized in this journal in Donaldson and Storeygard (2016). Furthermore, databases that assemble information from various sources provide disaggregated information on economic activity at a granular geographic level, such as the G-econ database (Nordhaus and Chen, 2006) that provides proxies of global income and population at the 1-arc degree.

in a network is nearby to other nodes with high eigenvector centralities. Here, it means that locations are more populated (and wealthier) the closer they are to other more populated (and wealthy) locations. Moreover, the eigenvalue of the system corresponding to this eigenvector turns out to be the welfare of the global economy (which is characterized by a single scalar because the infinitely elastic labor supply ensures welfare is equalized across all locations). In the more general case, the equilibrium of the spatial economy constitutes a network system of nonlinear equations. The properties of such systems remains an active field of research: Allen, Arkolakis, and Li (2020) offers a starting point.

We furthermore assume that all income accrues to labor, which allows us to recover average wages for a location given knowledge of income and population. This strong assumption clearly abstracts from other sources of income like capital, landholdings, firm profits, and others. One could argue that all these sources of income eventually accrue to individuals as well; indeed, as long as the income remains in a particular location, the predictions of the global spatial framework does not change by incorporating these other sources of income. (For example, as long as individuals in a location own their own homes, a model where individuals spend money on housing is no different – we say it is "isomorphic" – to the framework described above.) But in reality, not all income earned in a location accrues to the labor in that location, and such spatial flows of income would present another linkage between locations that we abstract from here.

Next consider data on economic linkages across space. As noted earlier, geographic distance is offers a convenient proxy for economic distance. But recently, researchers have begun to improve upon the distance proxy with measures of actual travel costs between locations. For example, Donaldson (2018) estimates the relative cost of traveling between locations via road, rail, and waterways by calculating the lowest cost route using Dijkstra (1959)'s algorithm – the same algorithm used by, for example, Google Maps. Allen and Arkolakis (2014) use a continuous space extension of the Dijkstra algorithm known as the Fast Marching Method (Tsitsiklis (1995); Sethian (1999)) to calculate travel times along the optimal route between locations. Allen and Arkolakis (2019) offers an analytical solution for the inverse economic distance as a function of the underlying transportation network.

Intuitively, these related approaches all share two advantages. First, they provide more precise estimates of the economic distance between two locations than distance alone would provide. (For example, Milwaukee and Grand Rapids are about 115 miles away as the crow flies, but travel between the two around Lake Michigan more than doubles the distance). Second, accounting for the underlying transportation network allows researchers to assess how changes in transportation infrastructure (e.g. improving the interstates I-90 and I-94 that connect Milwaukee and Grand Rapids) affects the spatial distribution of economic activity.

For any observed measure(s) of the economic linkages, the inverse economic distance $\{T_{ij}\}$ can then be constructed by regressing the observed (log) value of trade flows on those measures, conditioning on the origin and destination fixed effects. The predicted values of this gravity-model regression (excluding the estimated fixed effects) are the implied inverse economic distance.[8] For example, if one uses travel times as a measure of economic linkages,

[8] An alternative procedure would be to calibrate the inverse economic distance to exactly match the observed bilateral trade flows by including the regression residual in its construction. Such a procedure

the inverse economic distance would be the product of travel time and its estimated coefficient from such a regression.

Estimating supply and demand

Given measures of income Y_i in each location and a measure of the strength of the linkages T_{ij} between locations, we can calculate the global geography of every location that is, the inward and outward market accesses MA_j^{in} and MA_i^{out}.[9] We provide an iterative algorithm for solving that nonlinear system of equations in the companion Matlab code.

Now let us return to our augmented supply and demand equations for the global case. We observe the left-hand-side price variable, the wage for each location w_i, and the right-hand-side quantity variable, the population L_i. We also observe the data needed to calculate the market access variables (MA_i^{in} and MA_i^{out}).

We would like to estimate the coefficients on the right hand side variables, which represent the local and global elasticities of supply and demand. In doing so, the residual terms will be equal to measures of local productivity and local amenities (that is, $\ln C_i^D$ and $\ln C_i^S$).[10] Or put another way, we would like to estimate a system of supply and demand curves where we observe data on equilibrium outcomes of price and quantity at different times, which poses problems that are all-too-well understood!

How do we go about estimating our supply and demand curves? It might perhaps be more informative to start with what not to do. Following in the footsteps of Baldwin and Taglioni (2006), let us award medals for different types of errors that can arise, ranking them from most to least obvious.

The Bronze Medal Error

One glaring mistake in estimating supply and demand equations and – our "bronze medal" error – would be to use ordinary least squares regression. This approach is clearly not appropriate due to familiar simultaneity issues: what appears in data on wages and workers is the intersections of supply and demand curves, which does not trace out the shape of either

– which is closely related to the "exact hat algebra" pioneered by Dekle, Eaton, and Kortum (2008) and discussed in Costinot and Rodríguez-Clare (2014) – can result in an over-fitting problem when conducting counterfactuals (Dingel and Tintelnot (2020)).

[9] Recovering the global geography from the observed income and economic distances is a well-behaved problem. One can show using tools from Allen, Arkolakis, and Li (2020) that there exists a unique (to-scale) inward and outward market accesses MA_j^{in} and MA_i^{out} that solve the equations for any set of incomes Y_i and inverse economic distances T_{ij}.

[10] This approach of recovering the underlying geography based on the supply and demand residuals is equivalent (but perhaps easier to digest) to an approach that directly inverts the equilibrium market clearing conditions, as in Allen and Arkolakis (2014) and Redding (2016).

a supply or a demand curve, but rather a series of movements in both of them. (To put it another way, because the right hand side population variable is determined in equilibrium from equating supply and demand, it will be correlated with both the productivity and amenity shifters.) As a result, the coefficient from such an ordinary least squares regression will not recover either the supply or demand elasticity.

One strategy for overcoming this bronze medal error would be to employ instrumental variables: for example, using variation in the amenity $\ln C_i^S$ as an instrument for the equilibrium population to estimate the labor demand elasticity and using variation in the productivity as an instrument for the equilibrium population to estimate the supply elasticity. Conceptually, this involves looking at a source of shifts in labor supply (in this case, local amenities) to trace out a labor demand curve, and a source of shifts in labor demand (in this case, changes in local productivity) to trace out a labor supply curve. As long as the chosen instrumental variation in the amenities and productivities are uncorrelated, this will yield consistent estimates of the demand and supply elasticities.

What are some examples of such instruments? One example comes from Glaeser and Gottlieb (2009), who argue that the advent of air conditioning improved the amenity of locations with warm climates. Under the assumption that the climate of a location is not also correlated with the change in the productivity of a location, the climate of a location can be used as an instrument for change in population to identify the demand elasticity (for example, Allen and Donaldson (2020)).

Conversely, Allen and Donaldson (2020), following Bustos, Caprettini, and Ponticelli (2016), argue that increased global demand for soy improved the productivity of locations particularly well-suited for the production of soy. Under the assumption that the potential yield of soy in a location (say, relative to its potential yield for corn) did not also change the amenity of a location, the potential relative yield of soy to corn can be used as an instrument to identify the supply elasticity. Of course, the climate or agroclimatic properties are likely correlated with myriad characteristics of a location, making it unlikely these assumptions hold when comparing wages and populations across locations in cross section at a point in time. As such, it is preferable to rely on panel variation, looking at changes in wages and populations across locations over time (or, equivalently, including location fixed effects in the estimation of the supply and demand equations).

The Silver Medal Error

Somewhat less obviously, our "silver medal" error would be to ignore the spatial linkages between locations and simply estimate supply and demand using the local supply and demand equations based on the Rosen-Roback model. However, doing so ignores the variation in inward and outward market access across locations, relegating that variation to the residual term.

The instrumental variable strategy just described to address simultaneity bias is insufficient to address this bias. To see this, suppose you are estimating the labor demand equation, while using an amenity shifter like the arrival of air conditioning as an instrumental variable for population. Even if that amenity shifter is uncorrelated with productivities, it will be correlated with outward market access, biasing the estimate of the demand elasticity. Indeed, the only situation where this bias does not arise is in the special case when all locations share the same market access (as in the local spatial equilibrium).[11]

Fortunately, avoiding this mistake is straightforward: from the discussion above, one can construct measures of inward and outward market access measures from readily available spatial economic data. Including these market access measures in the supply and demand equations is a simple remedy to avoid the silver medal error.

The Gold Medal Error

An even more subtle concern is that outward and inward market access measures are themselves almost surely correlated with the productivity and amenity of a location. After all, the market access of a location depends in part on its own economic activity, which of course depends in equilibrium on its productivity and amenity. As a result, just including the market access measures in the supply and demand equations as controls will result not only in biased estimates of both the local and the global elasticities of supply and demand.

To address this concern, one can again use an instrumental variables strategy, instrumenting for both the population in a location and for the market access of that location. We discussed above possible instruments for the population; what about for market access? An appropriate instrument would be correlated with market access, but uncorrelated with local productivities or amenities.

In conceptual terms, think of market access as a type of inverse economic distance-weighted average of economic activity near a location. For an appropriate instrumental variable, suppose you use the measures of local productivities and amenities along with plausible values of the model elasticities to calculate the local equilibrium of a hypothetical economy using the basic local-area Rosen and Roback supply and demand equations. In

[11] Our "silver medal" error is similar in spirit to Baldwin and Taglioni (2006)'s "gold medal" error of failing to control for variation in market access in gravity equations. The two errors are distinct because unlike a gravity regression, the supply and demand regressions are not estimated using bilateral flows. As a result, their proposed solution of controlling for market access with origin and destination fixed effects does not apply here.

this hypothetical economy, spatial linkages do not matter and the only heterogeneity in productivities and amenities across locations arise from observables. Next, combine the implied equilibrium income in each location from this hypothetical economy with the observed economic distance and use the market access expressions above to calculate what the market access would be in such a hypothetical economy. This hypothetical market access measures how well connected each location is to the rest of the world, if the income in each location depended only on its observed productivities and amenities.

The hypothetical market access is a valid instrument for the actual market access under the assumption that observed productivities and amenities elsewhere in the world are uncorrelated with a location's own unobserved productivities and amenities. Using the hypothetical market access as an instruments then isolates the impact of market access on the supply and demand curves using this variation in productivities and amenities elsewhere through the spatial structure of the model.[12] Examples of such "model implied" instruments can be found in Monte, Redding, and Rossi-Hansberg (2018), Allen, Arkolakis, and Takahashi (2020), and Adao, Arkolakis, and Esposito (2021).

Taking Stock

Suppose you have successfully avoided the bronze, silver, and gold medal errors by estimating the labor supply and demand curves while appropriately using instrumental variables for the observed population and the market access terms. Now what?

You are now armed with estimates of the model elasticities, data on wages, populations, and market access terms, and with residuals terms from the supply and demand equations that correspond to the productivities and amenities in each location. Put another way, if you know the supply and demand elasticities, you can always find the local geography such that the observed distribution of economic activity – combined with the inverse economic distances you have constructed – is the global spatial equilibrium of the model.

Because you have recovered the geography that is consistent with the observed economic activity and you know the model elasticities, you are now able to assess how changes to the geography will affect the global spatial equilibrium. In the next section, we will discuss ways in which this approach can inform understanding the effects of various events and policy decisions.

[12] Another possibility would be to construct an instrument based on the augmented global supply and demand equations but excluding the own location (and perhaps also nearby locations) from the sum. But even if there is no spatial correlation in the productivity and amenity of locations, the equilibrium economic activity elsewhere depend in part on the economic activity of the own location (and hence the own productivity and amenity shifters), so such an instrument is unlikely to satisfy the exclusion restrictions.

5 Understanding the spatial impact of economic policies

We have seen how the global and local geographies interact through supply and demand to shape the spatial equilibrium and how those supply and demand curves can be combined with spatial data to apply the framework to the real world. Now we are equipped to describe the many types of questions that can be addressed with such a framework. We classify these questions into three types: those examining the impact of changes to the local geography, those examining the impact of changes to the global geography, and those which extend the framework above to incorporate additional spatial linkages beyond the flow of goods. We make no pretense here of offering a full survey of the literature; instead, our goal is to illustrate the extraordinary range of this work across events, policies, places, and times.

Local Geography Shocks

Consider first the question of how changes to local geography – changes to amenities which shift the supply curve or changes to productivities which shift the demand curve – affect the spatial distribution of economic activity.

Changes in the natural environment due to climate change offer many such examples. Rising sea levels and the resultant flooding both reduce the amount of land available for production and reduce the attractiveness of living in a coastal location, shifting both supply and demand curves in such locations inward, inducing populations to migrate elsewhere. Desmet, Kopp, Kulp, Nagy, Oppenheimer, Rossi-Hansberg, and Strauss (2018) studies the long-run impact of coastal flooding using a dynamic variation of the framework described here, finding that approximately 1.5% of the world population will be displaced by the year 2200 under current projections of the extent of flooding. Changing temperatures and patterns of precipitation also affect the suitability of different locations for producing different types of crops, affecting the productivity of different locations. Costinot, Donaldson, and Smith (2016) examine the long-run impact of estimated future changes in agricultural productivity across the globe to assess its impact on the spatial distribution of economic activity, estimating that climate change will result in a decline in the global value of agricultural output by approximately one-sixth.

Conflict and war can also reduce local productivities and amenities, although it remains an outstanding question for how long after the conflict these effects persist. For example, Davis and Weinstein (2002) examine the rebuilding of Japan after World War II, finding that the post-war distribution of economic activity closely mirrored the pre-war distribution, suggesting that war-time destruction was not enough to overcome fundamental characteristics different locations. In contrast, Chiovelli, Michalopoulos, and Papaioannou (2018) find the removal of landmines in the period after Mozambique's civil war had substantial impacts on the spatial distribution of economic activity, especially after accounting for the impacts of the demining on market access, i.e. the global geography.

Technological innovations may also increase the productivities in certain locations, shifting the labor demand curve outward. For example, Bustos, Caprettini, and Ponticelli (2016) presents evidence that the introduction of genetically modified soybeans in Brazil had heterogeneous effects on across areas with different soil and weather characteristics, and also was a labor-saving technology that ended up boosting industry. Caliendo, Parro, Rossi-Hansberg, and Sarte (2018) extend the framework above to incorporate intersectoral linkages along with spatial linkages to examine for example how local productivity improvements resulting from California's computer industry boom and the introduction of shale oil production in North Dakota affected the spatial distribution of economic activity. Some interesting topics for future research along these lines include the spatial effects of automation (as in Acemoglu and Restrepo (2020)) or new technologies that allow for remote work (as in Dingel and Neiman (2020); Althoff, Eckert, Ganapati, and Walsh (2022)).

Place-based policies enacted by the government can also be viewed as examples of shifts to the local demand or supply curves (depending on the particular nature of the policy). For example, Diamond and McQuade (2019) show that tax credits for low-income housing projects across 129 counties nationwide raised housing prices and reduced crime rates in low-income neighborhoods, but reduced housing prices in high-income neighborhoods. Some recent work seeks to characterize the trade-offs of such policies: for example, how policies that attract high-skill workers to low-wage cities can have broader social benefits and the equity-efficiency trade-offs of focusing place-based policies on locations with a dense concentration of low-income households (for discussion, see Fajgelbaum and Gaubert (2018); Gaubert, Kline, and Yagan (2021)).

Global Geography Shocks

Now let us turn our attention to how changes to global geography – changes in the economic distances and the resulting changes in the market access – affect the spatial distribution of economic activity.

Investment in transportation infrastructure which reduce the economic distance between locations is a natural application for evaluating changes global geography. For example, the US interstate highway system increased US welfare by 1-1.4 percent of GDP, more than its costs (Allen and Arkolakis (2014)); the US railroad system constructed in the second half of the 19th century more than doubled the price of land in nearby agricultural counties (Donaldson and Hornbeck (2016)); the Los Angeles Metro rail system increased commuting, but with little effect on productivity or amenities, and thus has considerably larger costs than benefits (Severen (2019)); the 1965, the Appalachian Development Highway System started in 1965 did benefit Appalachian counties, but most of the benefits accrued outside the region (Jaworski and Kitchens, 2019); and the arrival of the steam railway in mid-19th century London led to a doubling of population and land prices, as well as a geographical separation of workplaces and residences (Heblich, Redding, and Sturm 2020). Recent work has also examined the distributional implications of such infrastructure investments, for example, transportation infrastructure investments in New York City from 1870 to 1940 seem to have caused greater racial sorting and disparities (Lee (2022)) and the recently constructed national highway system in China benefits the economy of larger regional cities at the expense of rural regions (Baum-Snow, Henderson, Turner, Zhang, and Brandt (2020)).

While the simple framework above abstracts from the possibility that the economic distances may depend in part on the amount of trade between two locations, Duranton and Turner (2011) demonstrate the empirical relevance of congestion by showing that neither additional roads nor mass transit seem to reduce congestion in US cities. Recent work has made substantial progress incorporating congestion into spatial frameworks like the one described above. For example, Fajgelbaum and Schaal (2020) study optimal transportation networks in the presence of traffic congestion. In applying their framework to European countries, they find that the desirable network depends on whether they focus on flows within countries, or flows between countries. In a similar spirit, Allen and Arkolakis (2019) develop a spatial framework that includes congestion, and apply it the US highway network and the Seattle road network. These types of frameworks could also be used to evaluate congestion tolls imposing tolls in specific areas of the cities, such as the London or Singaporean traffic toll system or the congestion price system suggested for downtown Manhattan.

Other recent work, some still in a working paper stage, has sought to consider congestion in the context of ports, sea routes, and supply chains. In particular, the Allen and Arkolakis (2019) spatial framework for transportation and congestion has been applied to study the impact of several recent events in global shipping on the distribution of economic activity. For example, the 2016 expansion of the Panama Canal expanded trade between pairs of countries using the canal by 9-10 percent, although the costs of the expansion were borne by Panama (Heiland, Moxnes, Ulltveit-Moe, and Zi, 2019); the expansion of container shipping and Chinese-financed development of seaports across Africa and Asia is leading to reallocations away from more expensive ports like Singapore (Ducruet, Juhász, Nagy, and Steinwender, 2020); and entrepots, defined as shipping hubs that serve an intermediate role between place of origin and destination, play a key role in holding down global shipping costs (Ganapati, Wong, and Ziv, 2020). Another branch of this work looks at intermodal shipping: for example, how the construction of expressways in China early in the 21st century boosted exports (Fan, Lu, and Luo, 2019) and how to identify the nodes between road, rail, and ports in the US economy that would provide the greatest gains from additional investment (Fuchs and Wong, 2022). An exciting new area of work builds on the approach of Brancaccio, Kalouptsidi, and Papageorgiou, 2020 that develop a model endogenous route choices of exporters, and endogenous transportation costs, to study the global bulk shipping that constitutes 80% of world trade and evaluate the effect of large infrastructure projects such as the expansion of the Panama canal. Conwell (2022) combines endogenous route choices and traffic to find that an optimal subsidy on minibus entry in Cape Town may particularly benefits low-skill workers on long routes.

A classic example of changes in global geography arises from changes in international trade policy, e.g. changes in tariffs. For example, Topalova (2010) examines the impact of the the 1991 Indian tariff reduction to measure the impact of trade liberalization on poverty and rural districts, in which production sectors more exposed to tariff declines, experienced slower decline in poverty and lower consumption growth. The recent escalation of tariff measures by large economies such as U.S.A. and China has generated a renewed interest on the impact of tariff increases on the spatial distribution of economic activity, following the influential work of Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2020), who find that the recent U.S.-China trade war reduced U.S. real income by $7.2 billion, with the benefits of tariffs concentrated in politically competitive counties.

A final set of questions can be thought of as how changes to the local geography in some locations affect the economy elsewhere through the global geography. For example, beginning with the influential work of Autor, Dorn, and Hanson (2012), there has been much work on how productivity increases in China have affected workers in the U.S. and elsewhere through spatial linkages. Autor, Dorn, and Hanson (2012) find that US labor markets that previously included import-competing manufacturing industries experienced job and economic losses from the "rise of China." Caliendo, Dvorkin, and Parro (2019) use a spatial framework like the one above (expanded to include multiple sectors) to conclude that while there was an overall loss of manufacturing jobs from the rise of China, the US economy as a whole benefited, albeit with considerable variation across sector-state labor markets. The increase in demand elsewhere for goods or services in a location provides an other example: for example, Faber and Gaubert (2019) show that increasing international demand for tourism in Mexico causes large and significant local economic gains, which are in part driven by positive spillovers on manufacturing. In contrast, Allen, Fuchs, Ganapati, Graziano, Madera, and Montoriol-Garriga (2021) find that increasing international demand for tourism in Barcelona reduces the welfare of many local residents by increasing prices and crowding out local consumption.

Alternative Spatial Linkages

The framework developed above focuses on spatial linkages between locations that arise through the trade of goods. But of course people interact across space in many ways, including commuting, migration, or even social and business personal networks (see for example Christakis and Fowler (2009)). Some recent advances have incorporated other types of interactions into spatial frameworks like the one developed here.

Following the seminal work of Ahlfeldt, Redding, Sturm, and Wolf (2015), which considered how the rise and fall of the Berlin wall affected the spatial distribution of economic activity in that city, a number of papers have examined the impact of spatial interactions that arise through commuting flows. For example, the Severen (2019) mentioned earlier separates the commuting effect of the Los Angeles Metro from productivity or amenity effects, while Zarate Vasquez (2022) find that extensions of subway lines in Mexico City lead to increased commuting and a shift from informal to formal jobs. Monte, Redding, and Rossi-Hansberg (2018); Allen, Arkolakis, and Li (2015) combine commuting and spatial linkages in a single model: the first study finds that communities which win a competition for location of large plants have greater benefits if they have a more open commuting network; the second considers optimal zoning policy and finds Chicago would benefit from having more residences downtown and more business activity in outlying neighborhoods.

A related literature incorporates spatial linkages arising through altered migration patterns, extending the framework above to a dynamic setting. While the steady state (or balanced growth path) of these models resemble the static framework above, they are also able to yield predictions on the time it takes the economy to adjust to changes in geography. For example, in a global model with realistic geography Desmet, Nagy, and Rossi-Hansberg (2018) examine different scenarios for migration, and how eliminating migration restrictions would triple global welfare. Allen, de Castro Dobbin, and Morten (2018) show that walls built along the US-Mexico border altered migration patterns between Mexican municipalities and US counties. Tombe and Zhu (2019) argue that declining costs of internal migration in China can account for one-third of the aggregate growth in China's labor productivity from 2001 to 2005. Peters, 2021 finds that the expulsion of ethnic Germans from eastern Europe after World War II, and their return to West Germany, increased aggregate income per capita by about 12% after 25 years. Finally, Kleinman, Liu, and Redding (2021) find that the interaction of migration and capital investment can help to explain why convergence of incomes between US states declined between 1965 to 2015.

Another spatial linkage garnering recent attention is the formation of production linkages across firms. For example, lower costs of searching for and creating linkage between heterogeneous buyers and sellers can drive down marginal costs, as Bernard, Moxnes, and Ulltveit-Moe (2018); Bernard, Moxnes, and Saito (2019) find in applying their models to improved flow of people in Japan and to Norwegian customs data. Yet another spatial linkage can be measured by taking advantage of new data sets to assess the role of knowledge diffusion. Using nationally representative smartphone data, Couture, Dingel, Green, and Handbury (2020) examine patterns of travel and communication, while using highly granular smartphone data, Atkin, Chen, and Popov (2022) find substantial returns to what are actually face-to-face interactions in Silicon Valley. Using Facebook data grouped by zip code (and thus anonymized) (Chetty, Jackson, Kuchler, Stroebel, Hendren, Fluegge, Gong, Gonzalez, Grondin, Jacob, et al., 2022a,b) look at personal connections across socioeconomic groups and within cliques to study associations with economic mobility and determinants of connectedness.

Related studies look at the effects of new information technologies, documenting how the spatial spread of information can affect the distribution of economic activity. For example, Steinwender (2018) finds that the introduction of the trans-Atlantic telegraph 1866 provide information that affected cotton prices and trade flows, with gains equivalent to 8 percent of export value. Allen (2014) shows that including information frictions can make sense of observed patterns of regional agricultural trade flows prices in the Philippines. And Akerman, Leuven, and Mogstad (2022) find that this improved access to information in makes trade patterns more sensitive to distance and economic size using broadband expansion in Norway.

Recent research has incorporated even more types of spatial linkages including electricity transmission (Arkolakis and Walsh, 2022), piped water (Coury, Kitagawa, Shertzer, and Turner, 2022), and natural gas pipelines (Bachmann, Baqaee, Bayer, Kuhn, Löschel, Moll, Peichl, Pittel, and Schularick, 2022). The possibilities of adding additional spatial linkages or combining multiple types (or multiple layers) of linkages seem limitless. Moreover, extending the framework to include such interactions brings more realism and helps to illuminate the many ways in which geography shapes the spatial economy.

6 Conclusion

This article served three purposes. First, it was meant as an introduction to the reader about how geography shapes the spatial distribution of economic activity. In the classic Rosen-Roback framework, the answer depends solely on the "local" geography of each location and the equilibrium spatial distribution can be determined through familiar analysis of supply and demand curves. The major innovation of the new generation of economic geography models is to incorporate the spatial linkages between locations – putting space into the spatial model. The equilibrium can continue to be understood using the same supply and demand curves, but appropriately augmented to incorporate the impacts of the "global" geography.

The second purpose was to guide the reader through the process of combining these spatial models with spatial data to understand how geography shapes the real world spatial economy. Detailed spatial data are now readily available and researchers can apply these data to the theory using the well understood process of estimating supply and demand curves. With spatial linkages between locations arise potential pitfalls in estimation, but we offer strategies for traversing such issues. The end result is the ability to recover the underlying local and global geography such that the theory and data exactly correspond, allowing a researcher the ability to assess the impacts of any change in geography on the real world spatial distribution of economic activity.

Finally, we demonstrate the power of this close marriage between theory and data by highlighting the many types of questions that can be addressed. The types of questions and topics that can be examined using the framework here spans an incredibly wide range of topics, spanning economic history, environmental, labor, public finance, urban, and international, to name a few. This is an exciting time to be working on spatial issues: we have a new set of tools applicable to many interesting questions, most of which have yet to be tackled.

References

Acemoglu, D., and P. Restrepo (2020): "Robots and jobs: Evidence from US labor markets," *Journal of Political Economy*, 128(6), 2188–2244.

Adao, R., C. Arkolakis, and F. Esposito (2021): "General Equilibrium Effects in Space: Theory and Measurement," *mimeo*.

Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015): "The economics of density: Evidence from the Berlin Wall," *Econometrica*, 83(6), 2127–2189.

Akerman, A., E. Leuven, and M. Mogstad (2022): "Information frictions, internet, and the relationship between distance and trade," *American Economic Journal: Applied Economics*, 14(1), 133–63.

Allen, T. (2014): "Information frictions in trade," *Econometrica*, 82(6), 2041–2083.

Allen, T., and C. Arkolakis (2014): "Trade and the Topography of the Spatial Economy," *The Quarterly Journal of Economics*.

——— (2019): "The welfare effects of transportation infrastructure improvements," Discussion paper, National Bureau of Economic Research.

Allen, T., C. Arkolakis, and X. Li (2015): "Optimal City Structure," *Yale University, mimeograph*.

——— (2020): "On the equilibrium properties of network models with heterogeneous agents," Discussion paper, National Bureau of Economic Research.

Allen, T., C. Arkolakis, and Y. Takahashi (2020): "Universal gravity," *Journal of Political Economy*, 128(2), 393–433.

Allen, T., C. de Castro Dobbin, and M. Morten (2018): "Border walls," Discussion paper, National Bureau of Economic Research.

Allen, T., and D. Donaldson (2020): "Persistence and path dependence in the spatial economy," Discussion paper, National Bureau of Economic Research.

Allen, T., S. Fuchs, S. Ganapati, A. Graziano, R. Madera, and J. Montoriol-Garriga (2021): "Urban Welfare: Tourism in Barcelona," *Unpublished manuscript*.

Althoff, L., F. Eckert, S. Ganapati, and C. Walsh (2022): "The geography of remote work," *Regional Science and Urban Economics*, 93, 103770.

Anderson, J. E. (2010): "The Gravity Model," *forth. Annual Review of Economics.*

Anderson, J. E., and E. Van Wincoop (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93(1), 170–192.

Arkolakis, C., and C. Walsh (2022): "Clean Growth," *University of Columbia, mimeo.*

Armington, P. S. (1969): "A Theory of Demand for Products Distinguished by Place of Production," *International Monetary Fund Staff Papers*, 16, 159–178.

Atkin, D., M. K. Chen, and A. Popov (2022): "The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley," Working Paper 30147, National Bureau of Economic Research.

Autor, D., D. Dorn, and G. H. Hanson (2012): "The china syndrome: Local labor market effects of import competition in the united states," Discussion paper, National Bureau of Economic Research.

Bachmann, R., D. Baqaee, C. Bayer, M. Kuhn, A. Löschel, B. Moll, A. Peichl, K. Pittel, and M. Schularick (2022): "What if? The economic effects for Germany of a stop of energy imports from Russia," *Wirtschaftsdienst*, 102, 251–255.

Baldwin, R., and D. Taglioni (2006): "Gravity for dummies and dummies for gravity equations," Discussion paper, National Bureau of Economic Research.

Baum-Snow, N., J. V. Henderson, M. A. Turner, Q. Zhang, and L. Brandt (2020): "Does investment in national highways help or hurt hinterland city growth?," *Journal of Urban Economics*, 115, 103124.

Bernard, A. B., A. Moxnes, and Y. U. Saito (2019): "Production networks, geography, and firm performance," *Journal of Political Economy*, 127(2), 639–688.

Bernard, A. B., A. Moxnes, and K. H. Ulltveit-Moe (2018): "Two-sided heterogeneity and trade," *Review of Economics and Statistics*, 100(3), 424–439.

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

Bustos, P., B. Caprettini, and J. Ponticelli (2016): "Agricultural productivity and structural transformation: Evidence from Brazil," *American Economic Review*, 106(6), 1320–65.

Caliendo, L., M. Dvorkin, and F. Parro (2019): "Trade and labor market dynamics: General equilibrium analysis of the china trade shock," *Econometrica*, 87(3), 741–835.

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2018): "The impact of regional and sectoral productivity changes on the US economy," *The Review of economic studies*, 85(4), 2042–2096.

CFS (2017): "2017 Commodity Flow Survey Public Use File," *Bureau of Transportation Statistics*.

Chaney, T. (2018): "The gravity equation in international trade: An explanation," *Journal of Political Economy*, 126(1), 150–177.

Chetty, R., M. O. Jackson, T. Kuchler, J. Stroebel, N. Hendren, R. B. Fluegge, S. Gong, F. Gonzalez, A. Grondin, M. Jacob, et al. (2022a): "Social capital I: measurement and associations with economic mobility," *Nature*, 608(7921), 108–121.

——— (2022b): "Social capital II: determinants of economic connectedness," *Nature*, 608(7921), 122–134.

Chiovelli, G., S. Michalopoulos, and E. Papaioannou (2018): "Landmines and spatial development," Discussion paper, National Bureau of Economic Research.

Christakis, N. A., and J. H. Fowler (2009): *Connected: The surprising power of our social networks and how they shape our lives*. Little, Brown Spark.

Conwell, L. (2022): "Are There Too Many Minibuses in Cape Town? Privatized Provision of Public Transit," Discussion paper.

Costinot, A., D. Donaldson, and C. Smith (2016): "Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world," *Journal of Political Economy*, 124(1), 205–248.

Costinot, A., and A. Rodríguez-Clare (2014): "Trade theory with numbers: Quantifying the consequences of globalization," in *Handbook of international economics*, vol. 4, pp. 197–261. Elsevier.

Coury, M., T. Kitagawa, A. Shertzer, and M. Turner (2022): "The Value of Piped Water and Sewers: Evidence from 19th Century Chicago," Discussion paper, National Bureau of Economic Research.

Couture, V., J. Dingel, A. Green, and J. Handbury (2020): "Quantifying Social Interactions Using Smartphone Data," *Unpublished manuscript.*

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

Dekle, R., J. Eaton, and S. Kortum (2008): "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, 55(3), 511–540.

Desmet, K., R. E. Kopp, S. A. Kulp, D. K. Nagy, M. Oppenheimer, E. Rossi-Hansberg, and B. H. Strauss (2018): "Evaluating the economic cost of coastal flooding," Discussion paper, National Bureau of Economic Research.

Desmet, K., D. K. Nagy, and E. Rossi-Hansberg (2018): "The geography of development," *Journal of Political Economy*, 126(3), 903–983.

Diamond, R., and T. McQuade (2019): "Who wants affordable housing in their backyard? An equilibrium analysis of low-income property development," *Journal of Political Economy*, 127(3), 1063–1117.

Dijkstra, E. W. (1959): "A note on two problems in connexion with graphs," *Numerische mathematik*, 1(1), 269–271.

Dingel, J. I., and B. Neiman (2020): "How many jobs can be done at home?," *Journal of Public Economics*, 189, 104235.

Dingel, J. I., and F. Tintelnot (2020): "Spatial economics for granular settings," Discussion paper, National Bureau of Economic Research.

Disdier, A.-C., and K. Head (2008): "The puzzling persistence of the distance effect on bilateral trade," *The Review of Economics and statistics*, 90(1), 37–48.

Donaldson, D. (2018): "Railroads of the Raj: Estimating the impact of transportation infrastructure," *American Economic Review*, 108(4-5), 899–934.

Donaldson, D., and R. Hornbeck (2016): "Railroads and American economic growth: A "market access" approach," *The Quarterly Journal of Economics*, 131(2), 799–858.

Donaldson, D., and A. Storeygard (2016): "The view from above: Applications of satellite data in economics," *Journal of Economic Perspectives*, 30(4), 171–98.

Ducruet, C., R. Juhász, D. K. Nagy, and C. Steinwender (2020): "All aboard: The effects of port development," Discussion paper, National Bureau of Economic Research.

Duranton, G., and M. A. Turner (2011): "The fundamental law of road congestion: Evidence from US cities," *American Economic Review*, 101(6), 2616–52.

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

Eckert, F., and M. Peters (2022): "Spatial structural change," Discussion paper, National Bureau of Economic Research.

Faber, B., and C. Gaubert (2019): "Tourism and economic development: Evidence from Mexico's coastline," *American Economic Review*, 109(6), 2245–93.

Fajgelbaum, P., and C. Gaubert (2018): "Optimal spatial policies, geography and sorting," Discussion paper, Working Paper.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020): "The return to protectionism," *The Quarterly Journal of Economics*, 135(1), 1–55.

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

Fan, J., Y. Lu, and W. Luo (2019): "Valuing domestic transport infrastructure: a view from the route choice of exporters," *The Review of Economics and Statistics*, pp. 1–46.

Fuchs, S., and W. F. Wong (2022): "Multimodal Transport Networks," Discussion paper, mimeo, University of Oregon.

Fujita, M., P. Krugman, and A. J. Venables (1999): *The Spatial Economy: Cities, Regions, and International Trade*. MIT Press, Boston, Massachussetts.

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

Gaubert, C., P. M. Kline, and D. Yagan (2021): "Place-based redistribution," Discussion paper, National Bureau of Economic Research.

Glaeser, E. L., and J. D. Gottlieb (2009): "The wealth of cities: Agglomeration economies and spatial equilibrium in the United States," *Journal of economic literature*, 47(4), 983–1028.

Head, K., and T. Mayer (2013): *Gravity equations: Workhorse, toolkit, and cookbook*. Centre for Economic Policy Research.

Heblich, S., S. J. Redding, and D. M. Sturm (2020): "The making of the modern metropolis: evidence from London," *The Quarterly Journal of Economics*, 135(4), 2059–2133.

Heiland, I., A. Moxnes, K. H. Ulltveit-Moe, and Y. Zi (2019): "Trade from space: Shipping networks and the global implications of local shocks," *Available at SSRN 3504623*.

Helpman, E. (1998): "The Size of Regions," *Topics in Public Economics. Theoretical and Applied Analysis*, pp. 33–54.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012): "Measuring economic growth from outer space," *American economic review*, 102(2), 994–1028.

Jaworski, T., and C. T. Kitchens (2019): "National policy for regional development: Historical evidence from Appalachian highways," *Review of Economics and Statistics*, 101(5), 777–790.

Kleinman, B., E. Liu, and S. J. Redding (2021): "Dynamic spatial general equilibrium," Discussion paper, National Bureau of Economic Research.

Kline, P., and E. Moretti (2014): "Local economic development, agglomeration economies and the big push: 100 years of evidence from the tennessee valley authority," *Quarterly Journal of Economics*, 129, 275–331.

Krugman, P. (1991): "Increasing Returns and Economic Geography," *The Journal of Political Economy*, 99(3), 483–499.

Lee, S. (2022): "When Cities Grow: Urban Planning and Segregation in the Prewar US," Discussion paper, mimeo, University of Michigan.

Manson, S. (2020): "IPUMS National Historical Geographic Information System: Version 15.0," .

McCallum, J. (1995): "National borders matter: Canada-US regional trade patterns," *The American Economic Review*, 85(3), 615–623.

Monte, F., S. J. Redding, and E. Rossi-Hansberg (2018): "Commuting, migration, and local employment elasticities," *American Economic Review*, 108(12), 3855–90.

Nordhaus, W., and X. Chen (2006): "Geographically based economic data (G-Econ)," http://gecon.yale.edu/.

Peters, M. (2021): "Market size and spatial growth-evidence from germanyâs post-war population expulsions," Discussion paper, National Bureau of Economic Research.

Redding, S., and A. J. Venables (2004): "Economic geography and international inequality," *Journal of international Economics*, 62(1), 53–82.

Redding, S. J. (2016): "Goods trade, factor mobility and welfare," *Journal of International Economics*, 101, 148–167.

Redding, S. J., and E. Rossi-Hansberg (2017): "Quantitative spatial economics," *Annual Review of Economics*, 9, 21–58.

Roback, J. (1982): "Wages, rents, and the quality of life," *The Journal of Political Economy*, pp. 1257–1278.

Rosen, S. (1979): "Wage-based indexes of urban quality of life," *Current issues in urban economics*, pp. 74–104.

Sethian, J. (1999): *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*, vol. 3. Cambridge University Press.

Severen, C. (2019): "Commuting, labor, and housing market effects of mass transportation: Welfare and identification," *The Review of Economics and Statistics*, pp. 1–99.

Steinwender, C. (2018): "Real effects of information frictions: When the states and the kingdom became united," *American Economic Review*, 108(3), 657–96.

Tombe, T., and X. Zhu (2019): "Trade, migration, and productivity: A quantitative analysis of china," *American Economic Review*, 109(5), 1843–72.

Topalova, P. (2010): "Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India," *American Economic Journal: Applied Economics*, 2(4), 1–41.

Tsitsiklis, J. N. (1995): "Efficient algorithms for globally optimal trajectories," *IEEE transactions on Automatic Control*, 40(9), 1528–1538.

Zarate Vasquez, R. D. (2022): "Spatial Misallocation,Informality, and Transit Improvements : Evidence from Mexico City," Policy Research Working Paper Series 9990, The World Bank.

Figure 1: A supply shock in the local spatial equilibrium

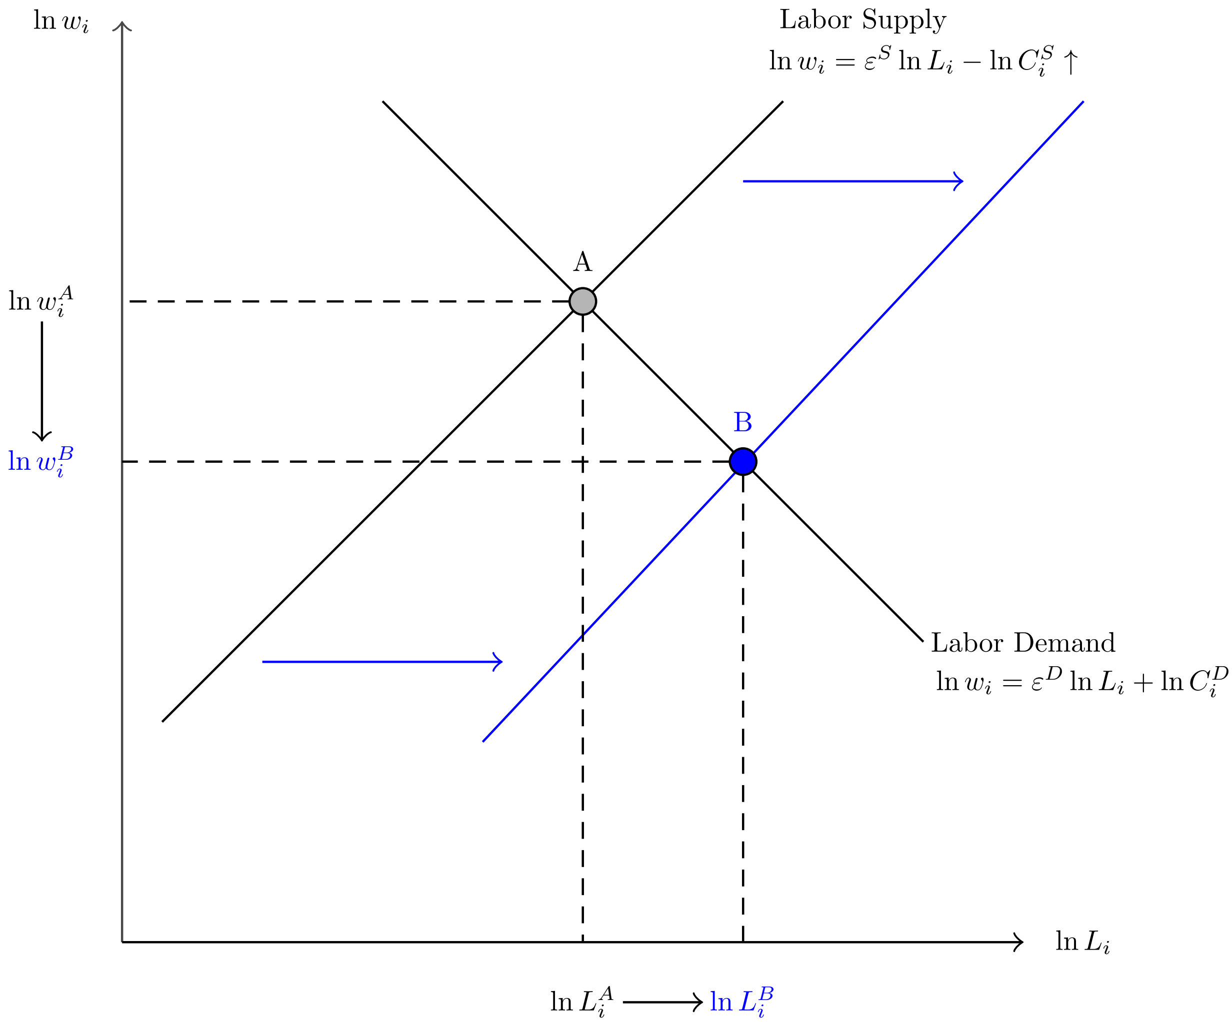

Notes: This figure illustrates the effect of an increase in the labor supply shifter on the equilibrium population and wages in a local spatial economy.

Figure 2: Spatial Linkages and Market Access

(a) Interstate trade flows

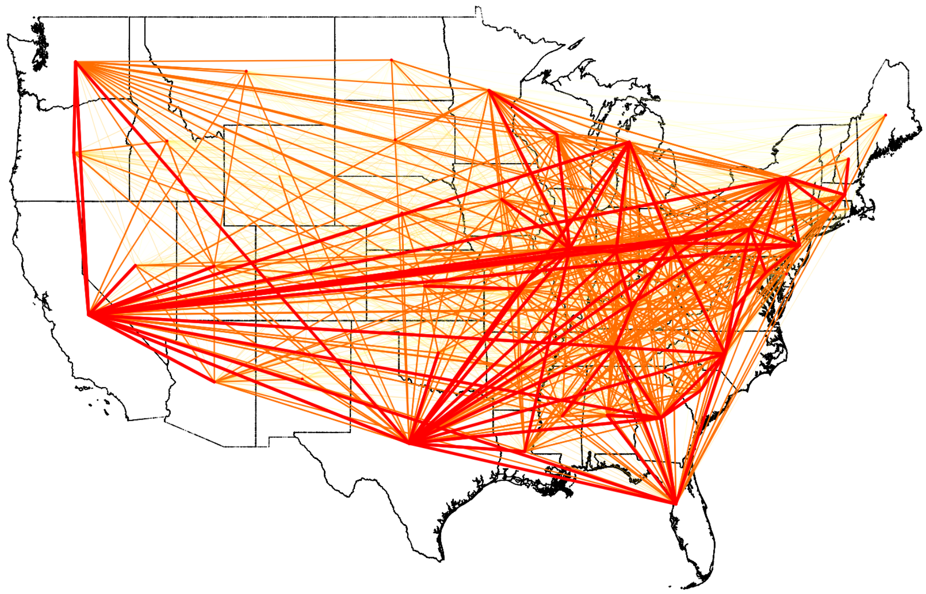

(b) Market access

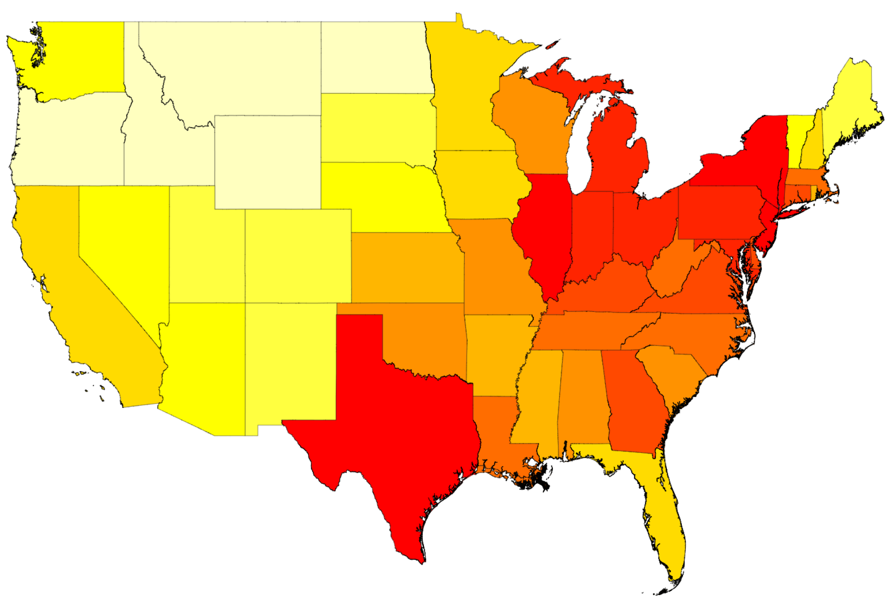

Notes: This figure illustrates the spatial linkages across U.S. states arising from trade flows. Panel (a) depicts the relative size of state-to-state bilateral trade flows, with thicker red lines indicating larger values and thinner yellow lines indicating smaller values. Panel (b) indicates the resulting (outward) market access of each state assuming trade costs T_{ij} are inversely proportional to distance.

Figure 3: A supply shock in the global spatial equilibrium

(a) The directly affected location

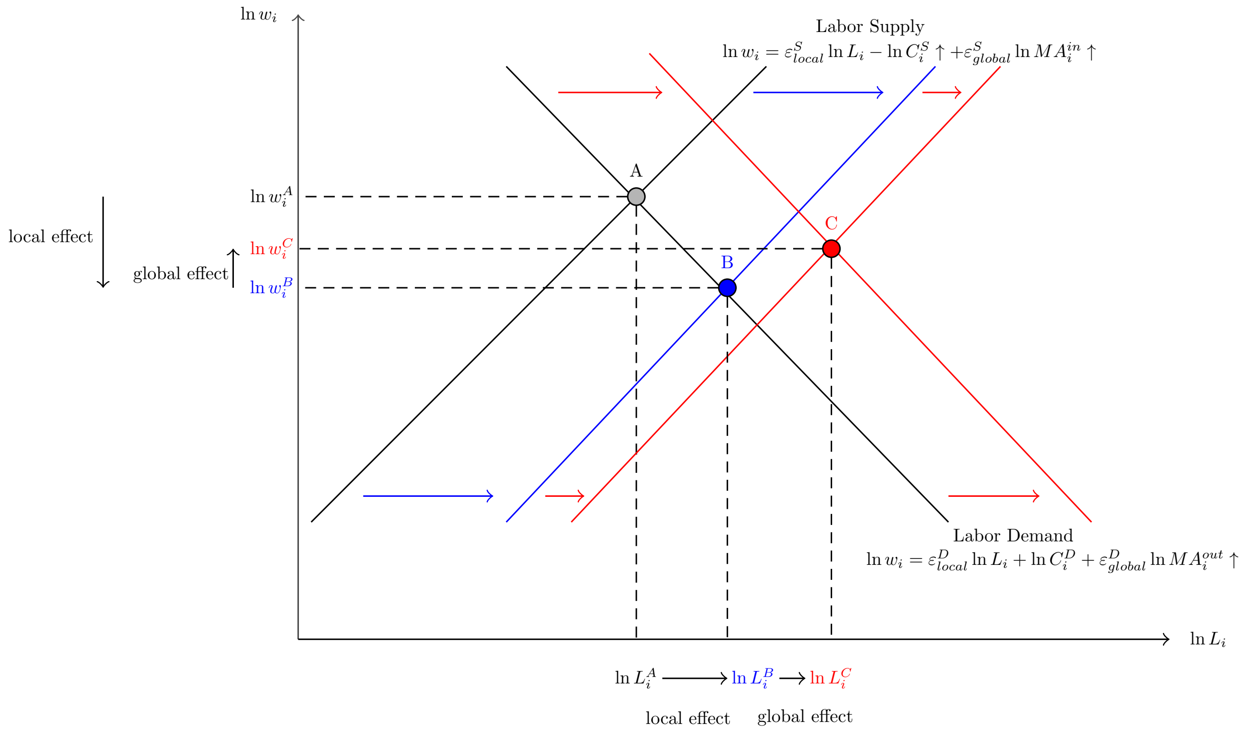

(b) An indirectly affected location

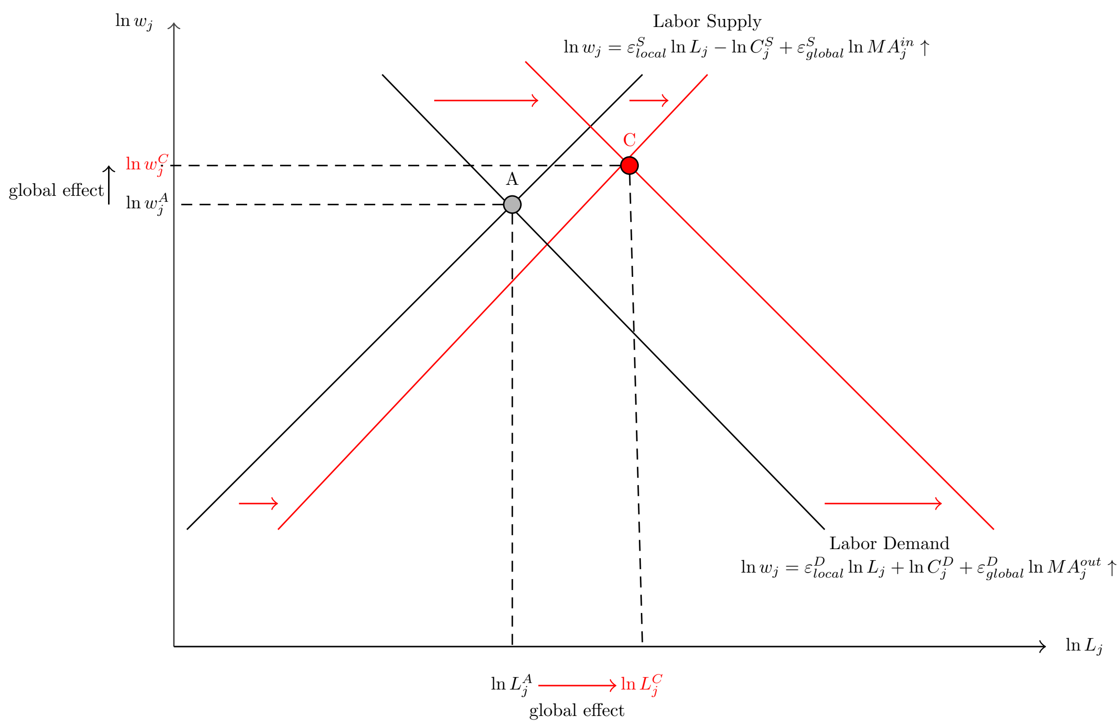

Notes: This figure illustrates the effect of an increase in the labor supply shifter in one location its own equilibrium population and wages (panel a) and another neighboring location (panel b).

A Appendix

In this Appendix, we provide detailed derivations of the results above.

A.1 A "Local" Spatial Economic Model

In this section, we provide a micro-foundation for the "local" supply and demand equations (1) and (2).

Consider a region comprising N different locations embedded in a larger economy. Agents can move freely across locations (or between these locations and the rest of the economy). Agents both produce and consume goods wherever they choose to live. Suppose that each location produces a homogeneous variety of good (e.g. corn). This good is freely traded across locations, but as all regions produce only that good there is no trade. We instead normalize the price of that good to one, $p = 1$, so that it provides a reference price to determine wages across locations.

The consumer problem is very simple in this context. Consumers maximize welfare,

$$W_i = c_i \times u_i \tag{7}$$

where c_i is the consumption in region i and u_i is a (non-consumption) amenity of residing in location i. The budget constraint of the consumer is simply $p \times c_i = w_i \iff c_i = w_i$, i.e. consumers consume an amount of the reference good equal to their wage.

The producer uses labor and capital to produce the final good. The production function in location i is given by:

$$Y_i = A_i K_i^{\theta} L_i^{1-\theta},$$

where $\theta \in [0, 1)$. Factor markets are perfectly competitive, so the marginal productivity of labor equals the wage and the marginal productivity of capital equals the rental rate:

$$w_i = (1-\theta) A_i K_i^{\theta} L_i^{-\theta}, \; r_i = \theta A_i K_i^{\theta-1} L_i^{1-\theta}.$$

It is evident that given capital the marginal productivity of labor declines with higher population and thus the wage in a location decreases when the population increases in that location. Assuming capital is fully mobile across locations so that the rental rate is equalized across locations (i.e. $r_i = r$), we have:

$$\frac{K_i}{L_i} = \left(\frac{r}{\theta A_i}\right)^{\frac{1}{\theta-1}}$$

and replacing in the wage equation yields:

$$w_i = (1-\theta)\, A_i^{-\frac{1}{\theta-1}} \left(\frac{r}{\theta}\right)^{\frac{\theta}{\theta-1}} \tag{8}$$

This model abstracts from a number of potentially important mechanisms, including other factors of production (like land), the consumption of non-tradables (like housing), possible heterogeneous preferences of different agents for different locations (e.g. I like beach and you like the mountains), economies of scale in production, etc. It turns out that a simple extension of the framework above is able to incorporate any combination of these different forces. Suppose that the productivity of a worker in a location depends in part on the total number of workers in that location:

$$A_i = \bar{A}_i L_i^{\alpha}, \tag{9}$$

where α may be positive or negative. Similarly, suppose that the amenity an agent derives from residing in a location depends in part on the total number of residents in that location:

$$u_i = \bar{u}_i L_i^{\beta}, \tag{10}$$

where again β may be positive or negative. In the model above, we have implicitly assumed $\alpha = \beta = 0$, but there are many reasons to think that α and β may be non-zero. For example, α may be negative if there is a fixed factor of production (like land) so that the more workers in a location, the less land there is per worker, driving down worker productivity. Alternatively, α may be positive if there are economies of scale in production. Similarly, β may be negative if residents also consume a local non-tradeable (like housing) that is in fixed supply, so that rent is driven up as the number of residents in a location increases. Or perhaps β is positive if greater population density induces greater supply of amenities (e.g. better parks). Allen and Arkolakis (2014) provide various micro-foundations for α and β along these lines.

Combining equations (8) and (9), we obtain the following labor demand equation:

$$\ln w_i = \frac{\alpha}{1-\theta} \ln L_i + \ln \left((1-\theta) \left(\frac{r}{\theta}\right)^{\frac{\theta}{\theta-1}} \left(\bar{A}_i\right)^{\frac{1}{1-\theta}} \right),$$

which is a special case of equation (1).

Similarly, because labor is perfectly mobile across locations, welfare equalization is equalized, i.e. $W = w_i \times u_i$. Combining welfare equalization with equation (10) yields the following labor supply equation:

$$\ln w_i = -\beta \ln L_i + \ln W \bar{u}_i^{-1},$$

which is a special case of equation (2).

A.2 A "Global" Spatial Economic Model

The goal in this section is to offer the derivations to the four equations comprising the equilibrium of the global spatial economy, namely equations (4), (4), (5) and (6). To do so, we rely on the same micro-economic foundations as in Allen and Arkolakis (2014), although as we discuss below, there are a number of alternative micro-economic foundations that also yield these equations (see e.g. Allen, Arkolakis, and Takahashi (2020)).

Consider a region comprising N different locations embedded in a larger economy. Agents can move freely across locations (or between these locations and the rest of the economy). Agents both produce and consume goods wherever they choose to live. Suppose that each location produces a distinct variety of good (e.g. French wine, Swiss cheese, ...). This assumption is called the "Armington" assumption (Armington (1969)). While clearly simplistic, it both makes the following derivations simpler and turns out to be mathematically equivalent to more realistic (but more complicated) models (see e.g. Eaton and Kortum (2002)).

Let us first consider the demand problem. Suppose that consumers like to consume many different varieties of goods. This "love of variety" creates an incentive for regions to trade with each other. In particular, we will assume that each agent has "constant elasticity of substitution" (CES) preferences such that if the agent lives in j and consumes quantity $\{q_{ij}\}_{j=1}^{N}$ of the variety of good from each location j she gets welfare:

$$W_j = \left(\sum_{i=1}^{N} q_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} u_j,$$

where $\sigma \geq 0$ is the elasticity of substitution (where higher σ indicates the agent is more willing to substitute one variety of good for another) and u_j is a (non-consumption) amenity of residing in location j. It is straightforward to show (but good practice to check!) that a utility-maximizing agent living in j with budget e_j and facing prices $\{p_{ij}\}_{i=1}^{N}$ will choose to spend $\{x_{ij}\}_{i=1}^{N}$, where:

$$x_{ij} = \frac{p_{ij}^{1-\sigma}}{\sum_{k=1}^{N} p_{kj}^{1-\sigma}} e_j.$$

and will receive welfare $W_j = e_j u_j / \left(\sum_{k=1}^{N} p_{kj}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$. Note that the share an agent spends of her income on each good does not depend on the level of her income, i.e. CES demand is homothetic. Since all agents living in a location face the same set of prices and CES demand is homothetic, we can can calculate the total amount spent by all agents living in location j on goods from i as:

$$X_{ij} = \frac{p_{ij}^{1-\sigma}}{\sum_{k=1}^{N} p_{kj}^{1-\sigma}} E_j, \tag{11}$$

where E_j is the total expenditure in j.

Now let us consider the supply problem. Suppose that the production of the differentiated varieties uses only labor as a factor of production and that each worker in location i can produce A_i units of the variety. Suppose too that it requires $t_{ij} \geq 0$ units of labor to ship a good from i to j. Finally, suppose that there is perfect competition, so that the price of goods are equal to their marginal cost of production. We then have that the price of a differentiated variety produced in i and sold in j is:

$$p_{ij} = \tau_{ij} \left(\frac{w_i}{A_i} \right), \tag{12}$$

where $\tau_{ij} \equiv 1 + t_{ij} A_i$ is the iceberg trade cost. Combining equations (11) and (12) yields the following gravity equation for trade flows:

$$X_{ij} = \frac{\tau_{ij}^{1-\sigma} \times \left(\frac{w_i}{A_i} \right)^{1-\sigma}}{\sum_k \tau_{kj}^{1-\sigma} \times \left(\frac{w_k}{A_k} \right)^{1-\sigma}} E_j. \tag{13}$$

Equation (13) is known as a *trade gravity equation* because it says that trade between locations is (a) proportional to the economic "size" of the origin and location; and (b) inversely proportional to the economic "distance" between the origin and destination. These two properties of trade flows are even more obvious when you re-write the equation as:

$$X_{ij} = T_{ij} \times \left(\frac{Y_i}{MA_i^{out}} \right) \times \left(\frac{E_j}{MA_j^{in}} \right), \tag{14}$$

where $T_{ij} \equiv \tau_{ij}^{1-\sigma}$ is the measure of (inverse) economic distance and we call $MA_j^{in} \equiv \sum_k \tau_{kj}^{1-\sigma} \times \left(\frac{w_k}{A_k} \right)^{1-\sigma}$ the *inward market access* and $MA_i^{out} \equiv \left(\frac{w_i}{A_i} \right)^{\sigma-1} Y_i$ the *outward market access*. As discussed in the main text, consumers in locations with higher inward market access benefit by being closer to the sellers of the goods they consume, whereas producers in locations with higher outward market access benefit from being closer to the buyers of the goods they produce. The new variant of the gravity equation in (14) highlights that the appropriate measure of economic size combines both the total income or expenditure of a location and its market access.

As noted in the main text, the inward and outward market accesses are closely related. To see this, we introduce two accounting identities. First, the income Y_i of each location i is equal to its total sales, i.e. $Y_i = \sum_{j=1}^{N} X_{ij}$, which when combined with the gravity equation (14) yields:

$$MA_i^{out} = \sum_{j=1}^{N} T_{ij} \times \left(\frac{E_j}{MA_j^{in}} \right). \tag{15}$$

Second, the expenditure E_j of each location j is equal to its total purchases, i.e. $E_j = \sum_{i=1}^{N} X_{ij}$, which when combined with the gravity equation (14) yields:

$$MA_j^{in} = \sum_{i} T_{ij} \times \left(\frac{Y_i}{MA_i^{out}} \right). \tag{16}$$

Equations (15) and (16) correspond to equations (3) and (4) in the main text. One neat thing about equations (15) and (16) is that given observed data on income and expenditures and estimates of (inverse) economic distances T_{ij}, you can use the two equations to uniquely identify (up-to-scale) the equilibrium inward and outward market access for every location. Another neat thing about the equations is that if the (inverse) economic distance is symmetric, i.e. $T_{ij} = T_{ji}$ and income is equal to expenditure, i.e. $Y_i = E_i$, then the inward and outward market access are equal up to scale, i.e. $MA_j^{in} \propto MA_j^{out}$ (which may be why oftentimes there is talk of "market access" without specifying if it is "inward" or "outward".).

Equations (15) and (16) let you calculate the inward and outward market access given information on income and expenditure. But how you figure out the equilibrium income and expenditure in each location? To close the model, we impose three market clearing conditions. The first market clearing conditions has to do with the demand for labor in a location. We require that the the income earned in a location is paid out to labor, i.e. $w_i L_i = Y_i$. This is straightforward in this model, as labor is the only factor of production and there is perfect competition, although the condition would have to be modified in models with multiple factors of production or with market power and firm profits.

Combining this equilibrium condition with the definition of outward market access (i.e. $MA_i^{out} \equiv \left(\frac{w_i}{A_i} \right)^{\sigma-1} Y_i$) yields the following labor demand equation:

$$\ln w_i = -\frac{1}{\sigma} \ln L_i + \frac{\sigma - 1}{\sigma} \ln A_i + \frac{1}{\sigma} \ln MA_i^{out}, \tag{17}$$

which is a special case of equation(5).

The second market clearing condition has to do with the supply for labor in a location. We assume that workers are equally happy to live in all locations and that level of happiness is in turn equal to the happiness they would achieve by living elsewhere in the economy. This comes from the assumption that workers are freely mobile across different locations: if workers can move wherever, why would anyone live in a location that makes them less happy? Of course, in reality, there may be many reasons that workers may live in locations with low levels of happiness, e.g. idiosyncratic preferences for different locations (more on this below) or the cost of moving between locations (which requires extending the static framework here into a dynamic one, see e.g. Desmet, Nagy, and Rossi-Hansberg (2018); Caliendo, Parro, Rossi-Hansberg, and Sarte (2018); Allen and Donaldson (2020)). Let us suppose that the level of happiness W_i an agent gets from residing in location i depends on both her utility from consumption and from a local amenity u_i. Finally, let us normalize the level of happiness in the rest of the world to one. While it may seems like a consequential choice to treat our set of locations as a small region in a large global economy, it actually is not: the equilibrium distribution of economic activity (i.e. the relative populations and incomes in all locations) is identical to a setting where the aggregate population is fixed.

Combining this equilibrium condition with the definition of inward market access (i.e. $MA_i^{in} \equiv \sum_k \tau_{ki}^{1-\sigma} \times \left(\frac{w_k}{A_k}\right)^{1-\sigma}$) yields the following labor supply equation:

$$W_i = 1 \iff \ln w_i = -\ln u_i - \frac{1}{\sigma - 1} \ln MA_i^{in}, \tag{18}$$

which is a special case of equation (6), albeit one where the labor supply is perfectly elastic.

Finally, we impose that income in equal to expenditure, i.e. $E_i = Y_i$. This implies that the value of goods being sent out of each location is equal to the value of goods being sent into each location, i.e. that trade is balanced. Since the model is a static one, this makes sense (although it highlights that the model is not well suited to explaining trade deficits observed in the data, which presumably arise due to dynamic considerations). Together, the labor demand equation (17), the labor supply equation (18), and the market access equations (15) and (16) can be solved together to determine the equilibrium population and wages in all locations.[13]

The model above provides an explanation for the market access equations (3) and (4), as well as special cases of the general labor supply and demand equations (5) and (6). As in the Rosen-Roback framework described in Appendix A.1, we can incorporate the presence of productivity and amenity spillovers of the form given in equations (9) and (10) to derive a more general form of the supply and demand curves.

Substituting equation (9) into the labor demand equation (17) yields the following mod-

[13]Because the equilibrium holds for any choice of units of wages, one also must choose a numeraire. In the companion Matlab code, we impose that the average wage across locations is equal to one.

ified labor demand curve:

$$\ln w_i = -\frac{1}{\sigma}\left(1-\alpha\left(\sigma-1\right)\right)\ln L_i + \frac{\sigma-1}{\sigma}\ln \bar{A}_i + \frac{1}{\sigma}\ln MA_i^{out}. \tag{19}$$

The more positive α, the flatter the downward sloping labor demand curve is, up to the point that $\alpha = \frac{1}{\sigma-1}$, at which point further increases in α actually cause the labor demand curve to shift upward! (It is at this point that a "black hole" equilibrium becomes possible where all population is concentrated in a single location, see Fujita, Krugman, and Venables (1999)).

Similarly, substituting equation (10) into the labor supply equation (18) yields the following modified labor supply curve:

$$\ln w_i = -\beta \ln L_i - \ln \bar{u}_i - \frac{1}{\sigma-1}\ln MA_i^{in}. \tag{20}$$

If β is negative (i.e. more people in a location reduce the amenity value of residing in that location), then the labor supply curve has a positive slope: to compensate perfectly mobile individuals for the amenity loss of the greater population, wages have to rise. As above, given the labor supply and demand equations along with the market access equations (15) and (16), we can solve the model to determine the equilibrium population and wages in all locations. But this has an interesting (and somewhat surprising conclusion): conditional on the slope of the labor supply and demand curves, the particular micro-foundation for a non-zero α and β do not matter for the equilibrium spatial distribution of economic activity. Or put another way, two different micro-foundations that both yield the same labor supply and demand curves will have the exact same implications for the equilibrium spatial distribution of economic activity.