Urban costs around the world

Jordan Rosenthal-Kay*
Federal Reserve Bank of San Francisco

March 27, 2026

Undergoing a substantial revision;
Please do not circulate this version

Abstract

Cities are engines of economic development, but the costs of urban scale limit the potential gains from urbanization. Urban costs depend on commuting costs and a city's capacity to expand up and out. Together, these determine a city's urban cost elasticity: the rate at which urban costs scale with population. Using a structural model and geospatial data on over 10,000 cities, I measure urban cost elasticities globally. Cities in the developing world grow by building out rather than up, even though residents face higher transportation costs, resulting in urban cost elasticities 35% higher than those in rich nations. Embedding my estimates of urban costs in a quantitative spatial model featuring a system of cities and rural-to-urban migration, I find that reducing urban cost elasticities globally to U.S. levels generates large welfare gains, both by drawing workers out of agriculture and by reallocating them toward previously inaccessible high-productivity cities. Turning to policy, I show that urban road paving is a cost-effective way to reduce urban costs in developing nations.

*Jordan.Rosenthal-Kay@sf.frb.org. I thank my advisors, Jonathan Dingel, Mikhail Golosov, Erik Hurst, and especially Esteban Rossi-Hansberg. This project greatly benefited from conversations with Rodrigo Adão, Milena Almagro, Elena Aguilar, Fernando Alvarez, Olivia Bordeu, Thomas Bourany, Juanma Castro-Vincenzi, Craig Chikis, Gilles Duranton, Ed Glaeser, Lu Han, Tom Hierons, Chang-Tai Hsieh, Greg Kaplan, Benny Kleinman, Hugo Lhuillier, John Mondragon, and Jeanne Sorin. I thank Ryden Iwamoto for excellent research assistance. All errors are my own.

1 Introduction

Cities around the world look different. Compared to cities in high-income nations, cities in the developing world have lower building heights, sprawl more, and are considerably denser, especially in their downtowns. In this paper, I argue that differences in urban form around the world reflect differences in *urban costs*: barriers to urban development. A city's urban costs depend on geophysical and regulatory constraints on development and the quality of its transportation infrastructure. Urban costs matter because, by limiting the size of cities, they limit an economy's ability to take advantage of the vast productivity benefits associated with urbanization. This paper asks, what barriers do cities around the world face to expanding, how do these barriers mediate the effects of urbanizing shocks, and what would the aggregate welfare gains be if these barriers were reduced? While substantial attention has been paid to quantifying the agglomeration benefits associated with urban life, this paper develops a framework to measure urban costs around the world and quantify their importance for national welfare.

Several facts suggest large economic gains from reallocating workers to the urban sector. First, there are large rural-urban wage gaps and substantial economic returns from rural-to-urban migration in the developing world (Gollin et al., 2014; Bryan et al., 2014).[1] Second, more populous and denser cities pay higher wages, controlling for selection on workers' unobserved abilities.[2] Urban costs constrain urbanization and therefore limit its benefits: if commuting costs or housing prices are high, or if the physical expansion of cities is burdensome, then it is difficult for people to agglomerate in space.

Understanding how urban costs affect economic development requires rich data on cities around the world. Traditional data sources on cities' wages, floorspace prices, and commuting patterns are not available with anything resembling worldwide coverage. Such data are particularly sparse in the developing world. To

[1] While these wage gaps may reflect the negative selection of low-productivity workers into agriculture necessary to relieve a nation's 'food problem' (Lagakos and Waugh, 2013) or the positive selection of high-productivity individuals to skill-intensive cities (Young, 2013; Behrens et al., 2014; Herrendorf and Schoellman, 2018), the literature nonetheless documents sizable returns for urban migrants (Lagakos et al., 2020). Gollin et al. (2021) provides suggestive evidence that urban wage premia in sub-Saharan Africa do not reflect compensating differentials for poor urban amenities. See Lagakos (2020) for a review.

[2] The literature finds estimates of wage elasticities to population and density around 3-5% globally. See Combes and Gobillon (2015) for a review of the literature estimating agglomeration elasticities. There is suggestive evidence that these elasticities may be higher in developing nations (Chauvin et al., 2017; Ahlfeldt and Pietrostefani, 2019). The literature has also paid considerable attention to the dynamic effect of moving to cities on workers' productivity. As sites of social interactions, cities foster human capital accumulation (Jacobs, 1969; Lucas, 1988; Crews, 2023), and empirical work documents steeper wage profiles for workers that move to cities (De La Roca and Puga, 2017; Lhuillier, 2023; Hong, 2024). The framework of this paper is static, but accounting for the dynamic gains of increased urbanization is an exciting avenue for future research.

overcome the challenge of data availability, I combine theory and satellite data on cities' urban form – their physical characteristics and spatial organization – to understand differences in urban costs around the world. Data on urban form are available for all cities. Moreover, these data are useful, because they allow me to infer urban costs from cities' spatial patterns of development: whether they have accommodated demand for floorspace by building vertically, horizontally, and relatively more in the urban core than the periphery.

I measure urban costs and quantify their importance in three steps. First, I develop a quantitative theory of cities that captures the key tradeoffs associated with urban scale and how they are reflected in urban form. Second, I use this theory to guide measurement of the parameters that determine the costs associated with scaling cities. Third, I embed my estimates in a general equilibrium model of an urban system which I calibrate to match data on over 10,000 cities in over 150 countries. I use the calibrated model to understand how urban costs affect a nation's level of economic development, and study cost-effective policies to improve cities.

To inform my theory of cities, I first define in the data what I mean by a 'city.' My definition begins with the entirety of the Urban Centres Database from the Global Human Settlement Layer, which identifies urban agglomerations defined by convex sets of remotely sensed built cover associated with a population of over 50,000 persons. These urban agglomerations can assemble large sets of cities (e.g., the San Francisco Bay Area, composed of San Francisco, Oakland, and San Jose) or large polycentric cities like Tokyo, Japan. I partition these agglomerations into monocentric 'cities' by developing an urban-theory informed estimator of the location and number of cities' central business districts. My estimator is not a routine application of off-the-shelf peak-finding or K-means type algorithms. Instead, it relies on the full spatial structure of cities' built environment, as measured by the World Settlement Footprint 3D (WSF-3D), a satellite observed (as opposed to modeled) dataset of built volume at a highly granular resolution (90m). I adapt tools from machine learning to tractably estimate these parameters for over 10,000 agglomerations and partition polycentric agglomerations into multiple cities. This procedure results in a dataset of over 15,000 'cities,' breaking up, for example, Tokyo into Yokohama, Chiba, and downtown Tokyo (Chiyoda).

Using my dataset of cities, I document facts on how urban form varies around the world. First, I find that polycentricity is predominantly a large-agglomeration, rich-world phenomenon: even conditional on city size, as measured by total agglomeration population, agglomerations in the rich world are more likely to be

polycentric. Thus, even adjusting for city size, cities in the developing world look more monocentric than their rich-world counterparts.

I then examine aggregate measures of urban form by examining the average height and density of cities around the world. Adjusting for city size, there is considerably less built volume per person in developing-world cities. This is predominantly driven by the fact that developing-world cities are substantially shorter than those in the rich world, and only modestly denser, indicating that cities in the developing world achieve their volume disproportionately with area rather than height: they build out, not up.[3] These results suggest that the technology to build cities varies around the world, and that there are either barriers to vertical development in the developing world, or horizontal barriers to development in the rich world.

Second, the internal structure of cities varies around the world. In developing cities, there is more built volume in the urban core relative to the periphery. I show this by first nonparametrically estimating the 'skyline slope' of cities around the world, which is the gradient of average height with respect to distance from a city's downtown. Skyline slopes are monotonically flatter in cities as one moves up the development spectrum, and surprisingly, log-linearity provides a strong approximation of the average skyline gradient. Consequently, I develop a parametric estimator of cities' skyline slopes and apply it to each of my fifteen thousand cities. This exercise reveals that even adjusting for city size, cities in the developing world have steeper skylines than rich-world cities. Through the lens of a simple urban model, this suggests either that commuting costs are greater in developing-world cities (concentrating demand at the urban core) or floorspace is more elastically supplied in the developing world, reducing the need for sprawl.

I use these features of the urban landscapes to measure urban costs. To do so, I develop a quantitative model of cities and an agricultural sector. In the model, differences in urban form reflect key parameters of the technologies that govern a city's capacity to expand. Cities crowd population at their core relative to their periphery if supply or demand for floorspace is greater in the city center. Supply is constrained by the capacity to build up, while demand depends on the transportation technology available in each city. Additionally, cities absorb population by building horizontally. Horizontal expansion offsets the costs of building up, but flattening cities increases the length of commutes and reduces the economies of density

[3] I define a city's average height as built volume divided by developable land area, so this contains not only the height of structures, but also the height of potentially developable land, like parks and parking lots.

that make cities productive.[4] All three of these aspects of the cities' building and transportation technology determine both the level of urban costs in each city, and each city's *urban cost elasticity*, which shapes how the costs of urban life scale with population size (Combes et al., 2019). My model incorporates key features from the urban and development literature. First, cities are monocentric and their internal structure depends on the transportation technology available in the city. Second, the economy is a dual-sector model, like that of Lewis (1954), with a traditional (agricultural) sector subject to decreasing returns and a modern (urban) sector subject to increasing returns. In contrast to Lewis (1954), in my model agricultural population is not in unlimited surplus nor is it perfectly elastically supplied to the urban sector. Finally, the model features a system of cities that vary in their technology, like Henderson (1974) or Desmet and Rossi-Hansberg (2013). Unlike these frameworks, in my model, cities are linked through costly trade and costless migration, following standard assumptions in quantitative economic geography (Allen and Arkolakis, 2014; Redding and Rossi-Hansberg, 2017).

The model generates several log-linear estimating equations that I use to recover three components of the cities' technology: a city's ability to build up (the floorspace supply elasticity), to build out (the land supply elasticity), and the commuting cost elasticity. First, a city's skyline slope reveals the product of the commuting cost and floorspace supply elasticities. The empirical challenge is to disentangle the two. The relationship between the average height of a city and its average income per square kilometer is informative of a city's floorspace supply elasticity. However, this relationship may reflect reverse causality: cities that are productive in building floorspace have lower floorspace prices, drawing in labor, making these cities larger and richer in equilibrium. Consequently, I instrument for a city's income per capita, controlling for population density, to identify the floorspace supply elasticity. I use model-generated instruments that recover the component of wages driven by a city's fundamental productivity, but not equilibrium prices nor endogenous productivity. Fundamental productivity is a valid instrument in this setting as it disentangles components of a city's wage that are driven by exogenous productivity in producing goods (for example, if it is resource-abundant) from general equilibrium effects driven by a city's size and its position in the intercity trade network. To recover land supply elasticities, which govern how responsive a city's footprint is to changes in land rents at its periphery, I use time series variation on urban growth in terms of area and income. Using the time series obviates the need for instruments as it allows me to control for city and region-year fixed effects that net out

[4] In the model, this channel operates by assuming that agglomeration forces operate through residential density.

the level and changes in horizontal building productivity. Cities in the developing world have both higher commuting costs and less elastic floorspace supply. Land supply elasticities are heterogeneous around the world and are primarily driven by differences in land availability.

I combine the estimated parameters to construct each city's urban cost elasticity, which is the elasticity of a city's consumption-equivalent utility to city population, holding wages and traded goods prices fixed. In my framework, the urban cost elasticity is a land-supply weighted average of congestion elasticities from rising floorspace prices and longer commuting costs. On average, cities in developing economies have urban cost elasticities that are 35% higher than those of the rich world. The largest driver of differences in urban cost elasticities across nations is differences in transportation technology, captured by the commuting cost elasticity.

I calibrate the model for nearly all countries to match their cities' population, output, building and transportation technology, and each country's intercity trade network and agricultural sector size and share of national income. To assess the macroeconomic impact of high urban costs, I first consider a counterfactual in which a nation's average urban cost elasticity is lowered to the average level observed in the United States, which is close to the global technological frontier. This yields large welfare gains around the world. On average, in the developing world, welfare rises by a staggering 66%, and the share of population living in cities rises by 8.5 percentage points. In the rich world, welfare rises by 11% and the size of the urban sector barely responds. In the developing world, around 20 percentage points of the welfare gain is driven by general equilibrium effects, as workers move out of agriculture and reallocate to more productive cities. This exercise underscores the potential scope for lower urban costs to foster economic development.

To study one type of policy that might achieve these gains, I focus on road paving. Many developing-world cities are substantially under-paved, and the costs of road paving are well understood. For example, I find that only 50% of urban roads in Nigeria are paved. The benefits of transportation infrastructure are difficult to monetize, which may explain why developing countries have under-invested in these technologies. I use data on cities' transportation technology to estimate the effect of road paving on my estimates of cities' commuting cost elasticity and find that a 10pp increase in the share of roads paved lowers the commuting cost elasticity by over 3%. To assess the efficacy of road paving as a policy, I implement a policy in which each country can spend at most 1% of baseline GDP on road paving. Countries use those funds to pave urban roads starting

with the most productive city first, and working down the city size distribution until funds are exhausted or all cities' roads are paved. The policy has little bite in the developed world, but it is cost-effective in many parts of the developing world, raising national welfare by well over 1%.

Related literature This paper relates broadly to three bodies of work. First, it contributes to the literature examining global urban patterns by providing new facts on how cities' internal structure varies globally. Second, it advances the literature that examines city-level policy interventions using quantitative spatial models by developing a framework to understand the aggregate and distributional implications of such changes across countries. Third, by studying the determinants of the overall size of the urban sector and the population distribution across cities, this paper contributes to the literature on how the spatial (mis-)allocation of productive factors affects macroeconomic aggregates.

Several recent papers have studied various aspects of cities globally using geospatial data. Akbar et al. (2023a) and Akbar et al. (2023b) document that there are large differences in urban speed around the world, primarily driven by differences in freeflow speed. They find that cities in the rich world are 50% faster than those in developing nations. Harari (2020) documents that cities' spatial layout affects their growth and amenity value in India. I provide a quantitative framework that links a city's characteristics like speed and urban form to its size. The studies most closely related to the empirical component of this paper are those by Jedwab et al. (2021) and Lall et al. (2021a), both of which study how cities' physical characteristics vary with city and national income. Jedwab et al. (2021) restrict their attention to around 1,000 large agglomerations and report relationships between urban income and city height and area similar to those presented in this paper. Lall et al. (2021b) and Lall et al. (2021a) also document relationships between urban form and city income.[5] The regressions reported in those papers lack a structural interpretation of the estimands and a consideration of potential endogeneity biases that confound the identification of model parameters. I additionally use information on cities' internal structure to understand differences in urban form globally.[6] I propose and quantify a theory of cities that explains variation in urban form as driven by differences in the building and transportation technology available around the world. I measure components of these technologies using

[5]There is also a large remote-sensing literature on the spatial patterns of urban development around the world. For example, Frolking et al. (2024) analyze how urban land use patterns have changed over time using satellite data, and find that Asian cities have shifted from building horizontally to building vertically.

[6]In a similar spirit, Mills and Tan (1980) collate population density gradient estimates for under 100 cities, and, similar to this paper's findings, finds steeper population density gradients in developing countries. Liotta et al. (2022) compute population density and rent gradients for 192 global cities to test predictions of the monocentric city model and find that the vast majority of cities globally have negative population density and rent gradients, consistent with theory.

model-consistent regressions, and I use my estimates to perform counterfactual analysis with a quantitative spatial model of the world. Two existing papers use frameworks similar to mine to answer related questions. First, Ahlfeldt et al. (2023) use data on the construction of tall buildings in cities globally to quantify the welfare gains with changes in building technology since 1975, while Coeurdacier et al. (2022) use data on changing agricultural and urban land-use patterns in France to understand how structural transformation has affected how cities grow. Unlike these papers, I focus on all aspects of urban costs: the ability to build up, out, and transportation technology, and apply it to all cities around the world.

Second, this paper examines the aggregate effects of city-level policy interventions using tools from quantitative spatial economics. Past research has used case studies to understand the effects of changes to transportation technology within cities. For example, the effect of introducing Bus Rapid Transit systems on urban residents' welfare has been studied in Jakarta, Indonesia (Gaduh et al., 2022; Kreindler et al., 2023), Bogotá, Colombia (Tsivanidis, 2023) and Dar es Salaam, Tanzania (Balboni et al., 2020).[7] Unlike these papers, this work moves beyond case studies and studies how city-level policy interventions matter for the macroeconomy by accounting for equilibrium interactions across all cities within a country. My framework provides less detailed treatment of city structure in order to be comprehensive and draw comparisons around the world. Moving beyond case studies meets the demands from policymakers. For example, the World Bank writes that in Indonesia, congestion forces 'limit the productivity gains of agglomeration' and Indonesia must '[a]ugment the coverage and quality of basic services and infrastructure for all people in all places' (Roberts et al., 2019). My framework is designed to assess the aggregate implications of high urban costs across all countries and quantify the potential welfare gains associated with reducing such costs.

Third, this paper is about the role of cities in determining national welfare, and the urban costs which limit their size. Most similar in spirit to this work are Bryan and Morten (2019), who develop a model of frictional interregional migration, which they use to estimate the aggregate productivity effects of reducing migration barriers in Indonesia, and Desmet et al. (2017) who study the productivity effects of relaxing migration barriers in Asia. Lagakos et al. (2023) study the welfare gains of rural-to-urban migration in a dynamic model in which seasonal migration occurs as a reaction to uninsured negative shocks faced by agricultural households. Unlike these papers, which focus on the barriers faced by agricultural migrants, I instead focus

[7] See Bryan et al. (2020) for a review of this literature.

on cities' urban costs, which limit their size.[8] Au and Henderson (2006) estimate agglomeration benefit and urban cost elasticities for cities in China, and find that Chinese cities are undersized relative to the social optimum. Desmet and Rossi-Hansberg (2013) study misallocation across cities. They examine how level differences across cities in the labor wedge contribute to labor misallocation across cities, and find that were these differences eliminated, there would be much larger gains to labor reallocation across cities in China than the United States. My framework allows me to study the macroeconomic impact of differences in cities' building and transportation technology across cities for the entire world.

The rest of this paper is organized as follows. First, in Section 2, I develop a dataset of cities around the world, partitioning polycentric agglomerations into monocentric cities. Using these data, I show facts about cities globally. I then present a quantitative model of cities aimed at rationalizing these facts in Section 3. In Section 4, I use model-implied estimating equations to recover urban costs for all cities around the world. I embed these estimates in my quantitative spatial model and perform counterfactual analysis in Section 5. Section 6 concludes.

2 Urban form around the world

This section presents facts on how cities differ globally in terms of their urban form: both aggregate characteristics – their area, volume, and density – and their internal spatial structure.

Doing so first requires that I define what a city is in the data. I begin by creating a dataset of urban agglomerations within which I can measure the built environment. I then combine theory and machine learning tools to partition these agglomerations into a dataset of cities (e.g., splitting the San Francisco Bay Area into San Francisco, Oakland, and San Jose).

[8] In my model, I allow for a finite elasticity of rural-to-urban migration with respect to changes in urban utility, but do not model bilateral migration costs, as I lack global bilateral data on rural-to-urban migration patterns.

2.1 Building a dataset of cities

To define the location and spatial extent of a city, I use the shapefiles provided by the Urban Centre Database (GHS-UCDB, R2024a), via Copernicus, a program of the European Commission. The UCDB provides 11,422 polygons associated with urban developments of more than 50,000 people. These polygons are defined by combining a variety of modeled and remotely sensed data products compiled by Copernicus, and correspond to the intuitive notion of an urban agglomeration: a contiguous collection of cells with a population density exceeding 1,500 persons per square kilometer. I subset the sample to agglomerations with a 'high' quality flag, resulting in a dataset of 10,852 agglomerations. Importantly, this definition is remotely sensed, and does not make reference to political boundaries in its construction.

Within each agglomeration, I measure contemporary construction using the World Settlement Footprint 3D (WSF-3D, Esch et al., 2020; Esch et al., 2022), which combines observations from the Sentinel-1 and -2 satellites with a digital elevation model (TanDEM-X) to produce rasters of built volume at the 90m resolution globally. These data are superior to products like the built volume data in the Global Human Settlement Layer, which combines modeled height data with remotely sensed built surface data to infer volume.[9] Instead, the WSF-3D more closely resembles satellite-observed volume than a machine-learned product. Moreover, its measurement is globally consistent, guaranteeing measurement error is orthogonal to the level of development and the availability of validation data.

Within each urban agglomeration, I extract the 2015 Visible Infrared Imaging Radiometer Suite (VIIRS) nighttime luminosity measure ($\approx$500m resolution) and population by way of WorldPop, which provides a lightly modeled population raster at the 100m level that does not use the built environment data as a modeling input.[10] Finally, I use a digital elevation model (Copernicus DEM GLO-30) combined with land cover data (classifications of land as soil, ice, and so on) to create a 10m-resolution raster indicating the potential for development to occur in that pixel, which I then aggregate to the level of the WSF-3D data. I classify land as developable if it has a slope with grade less than 15% and is not water, ice, wetland, or a mangrove.

[9] Previous versions of this project relied on the GHSL data.

[10] That is, I use the unconstrained data, which downscales administrative population estimates (census and projections) using covariates like land cover (water, ice, etc).

2.2 Identifying polycentricity: urban-theoretic partitions

The urban agglomerations in the data reflect an assemblage of what one might intuitively call 'cities.' For example, the San Francisco Bay Area polygon contains San Francisco, San Jose, and Oakland. Other agglomerations like Istanbul, Turkey, are famously polycentric and contain, at minimum, two 'cities' defined by the 'European' and 'Asian' sides, split by the Bosporus. The analysis throughout this paper defines cities as monocentric, convex clusters that may partition these larger urban units. In this section, I explain how I partition urban agglomerations into (potentially) smaller monocentric cities.

Theoretically, monocentricity means that there is a single central business district (CBD) to which agents commute for work, and the rest of urban land is used to house commuters to the CBD. I treat the number and location of these CBDs as unknown parameters to be estimated.

Suppose cities $k = 1, ..., K$ are circular with locations indexed in polar coordinates, (x, ϕ), where x denotes distance to a CBD, and ϕ is the polar angle.[11] Households in the city have a utility function, $U\left(C, A(x)h(x)\right)$ which they maximize by choosing C, $h(x)$ and x, subject to some budget constraint, $C + r(x)h(x) \leq w$. C is some non-floorspace ($h(x)$) aggregator and w is (for now) treated as exogenous. Their location decision x trades off amenities $A(x)$ and floorspace prices $r(x)$. Amenities decay with respect to distance to the downtown, $A(x) \propto x^{-\tau}$, where τ captures commuting costs, which are paid for in time not spent at home enjoying $h(x)$.[12]

Developers supply floorspace with a constant returns technology in land (a fixed factor) and materials, resulting in upward sloping housing supply curves $h(x) = \bar{h}(x)r(x)^{\gamma}$, where $\bar{h}(x)$ is some exogenous shifter of the quality or quantity of land.

Then, in equilibrium, we have that the total amount of floorspace at location i in city k is

$$H_{ik} = \bar{H}_{ik}\left(x_{ik}\right)^{-\theta_k} \tag{1}$$

[11] That is, space is continuous. Have no fear, reader: I only consider symmetric allocations of economic activity. Hereon we can ignore ϕ; its introduction is meant to underscore that cities have some nonzero area.

[12] An alternative is that commuting costs are paid for in units of labor. While standard – this assumption has the unattractive property that purchasing power for housing declines across space, which generates variable elasticity price gradients that are sensitive to the income and price elasticity of floorspace demand. Mills (1972) failed to account for this fact, and this feature of the monocentric city model was not corrected until Brueckner (1982).

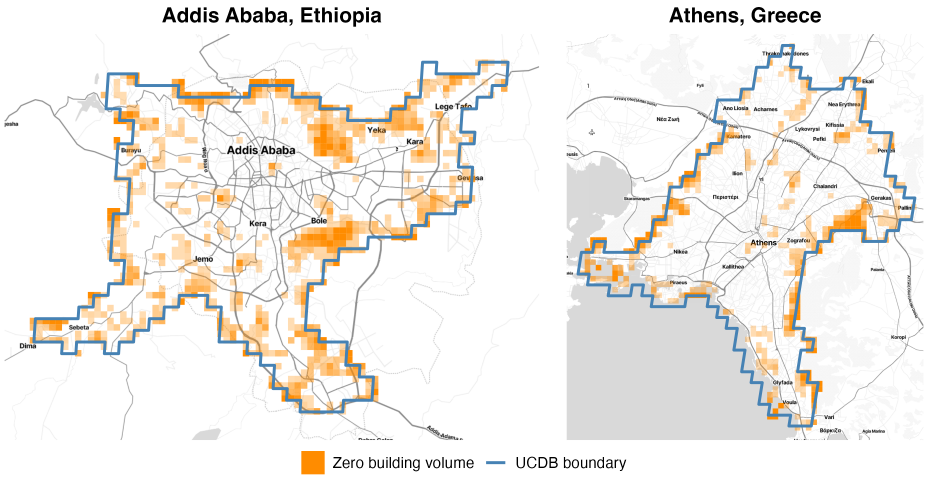


Figure 1: Structural zeros in Addis Ababa, Ethiopia, and Athens, Greece UCDB agglomerations. Addis is around 465km^2 with a population of 7.3M while Athens is 412km^2 with a population of around 3.2M. Colored pixels are aggregated to the 500m resolution and shaded by the share of zero-value WSF-3D pixels.

where $\theta_k = \tau_k \gamma_k$.

This model serves as the basis for CBD estimation and city partition. I observe building volume $H_i \geq 0$ on a grid of N spatial blocks within each urban agglomeration, as well as the fraction of land in each pixel that is developable, f_i. I rescale each urban agglomeration in the UCDB data onto the unit square $[0,1]^2$ and assign each pixel a location in the unit square corresponding to its centroid, μ_i.

I assume the statistical model,

$$H_i \sim \text{Poisson}\left(f_i \cdot \sum_{k=1}^{K} A_k (x_{ik})^{-\theta_k} \right) \tag{2}$$

where pixel i's distance to the CBD k, x_{ik}, is Euclidean. The key assumption here is additivity: that pixels aggregate over unobserved partitions of space in which agents are living in separate buildings. The rationale here is the superposition property of Poisson random variables, that if $H_{ik} \sim \text{Poisson}(A_k (x_{ik})^{-\theta_k})$, then $H_i = \sum_k H_{ik} \sim \text{Poisson}(\sum_{k=1}^{K} A_k (x_{ik})^{-\theta_k})$. As space is continuous in the model, the correct model-based analog of H_i is height, not volume. While pixels are of constant area, their developable area differs; f_i is the offset that corrects the scaling of volume data so that it is in constant proportion to height.

I use the Poisson model to rationalize zeros in the data that are plausibly structural: land with no development

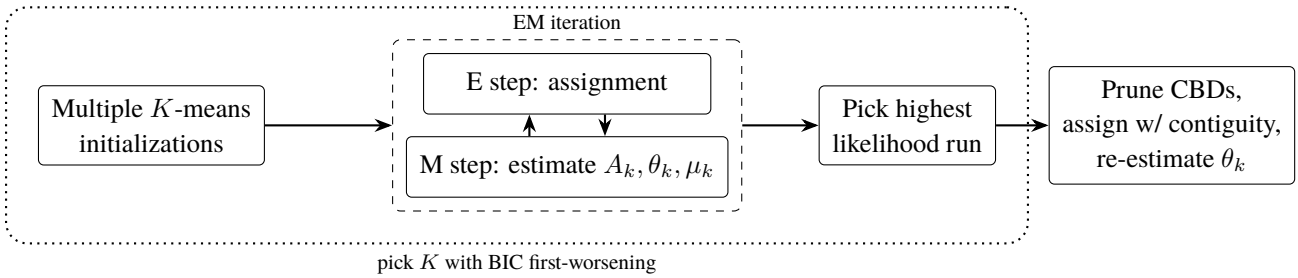


Figure 2: CBD detection algorithm for estimating the parameters Ω and K in (3).

is more likely to occur on the fringe of a city, rather than near its CBD. Thus, accounting for these zeros will meaningfully affect estimates of θ_k. Upon inspection of the data, these structural zeros are far more common in cities' peripheries than their downtowns, and more so in developing cities; see Figure 1 for an example of Athens, Greece and Addis Ababa, Ethiopia, which has about 4pp more empty pixels within the UCDB boundary.

The log likelihood for (2) is,

$$\ell(\Omega) = \sum_{i=1}^{N} H_i \ln \lambda_i - \lambda_i, \quad \lambda_i = f_i \sum_{k=1}^{K} A_k (x_{ik})^{-\theta_k}, \quad \Omega = \{\mu_k, \theta_k, A_k\}_{k=1}^{K}, \tag{3}$$

where $\mu_k \in [0,1]^2$ are the unknown CBDs. Solving for the unknown CBDs resembles the problem of K-means clustering and estimation is related, but it is not an off-the-shelf application for two reasons. First, the location of the CBDs must be jointly estimated with θ_k and A_k. Second, unlike K-means, in which space does not matter within a cluster assignment, here I care explicitly about spatial structure: I am looking for an isometric decay pattern centered at μ_k.

Estimating Ω by maximizing (3) is therefore nontrivial. I do so with an expectation-maximization (EM) algorithm. The EM algorithm consists of an iterative two-step process that alternates between an 'E' step and an 'M' step. The E step attributes mass in each cell i to each CBD k given the parameter estimates Ω. The Poisson model makes this attribution trivial: The expected contribution of H_i to CDB k is $\hat{H}_{ik} = a_{ik} H_i$ where $a_{ik} \propto A_k (x_{ik})^{-\theta_k}$. Given the assignment weights a_{ik}, I proceed to the M step in which A_k, θ_k and μ_k are updated. In particular, θ_k solves the weighted Poisson first-order condition and A_k can then be backed

out in closed form.[13] To update μ_k, I take Newton steps by differentiating the likelihood with respect to μ_k at the estimated parameter values and assignments. This is possible as μ_k enters through the distance function that produces $x_{ik} = \|\mu_i - \mu_k\|$.

I iterate between the E step and the M step until the likelihood function converges. In practice, I make a number of modifications to the algorithm to improve functionality: first, I aggregate data to the 500m resolution which massively reduces compute time, making it computationally feasible to estimate for every agglomeration in the data. I ensure that CBDs cannot be placed on undevelopable land with a guard on the Newton step, and place a soft penalty on the Newton step that prevents CBDs from colliding, ensuring separation of clusters as the algorithm explores the space. Additionally, I place a penalty on estimating $\theta_k < 0$ only during CBD detection to ensure the algorithm finds 'peaks' rather than 'troughs' in the spatial distribution of built volume.

I initialize guesses of μ_k using K-means, then use multiple restarts of the EM algorithm with different initializations and scale the number of restarts with city size. I choose $\hat{\Omega}$ from the highest-value likelihood among the set of restarts.

I select K using Bayesian information criterion (BIC), a principled way of penalizing additional parameters that do not sufficiently increase the likelihood. I estimate CBDs for $K = 1, 2, ...$ with a first-worsening rule: if BIC(K) <BIC$(K-1)$ I set the city's number of centers $K^* = K - 1$.

With K^* and $\hat{\Omega}$ in hand, I first prune CBDs with near-zero $\hat{A}_k$ or $\hat{\theta}_k \leq 0$, or for whom $\sum_k a_{ik}$ amounts to less than 2% of the total area.[14] I then assign pixels to a single CBD using a Potts model, which finds a hard assignment of pixels that maximizes the likelihood conditional on contiguity constraints. When assignments result in disconnected sets, I keep only the cluster that contains the CBD and reassign disconnected territories to the next-best CBD.

Once CBDs have been detected, I reestimate θ_k in a second step using the full, spatially granular WSF-3D data. In this second step, I form larger grids by binning up polar angles and radii, and weight observations

[13] That is, θ_k solves,

$$\sum_i \left(\hat{H}_{ik} - f_i A_k (x_{ik})^{-\theta_k}\right) \ln(x_{ik}) = 0, \quad A_k = \frac{\sum_i \hat{H}_{ik}}{\sum_i f_i (x_{ik})^{-\theta_k}}.$$

[14] I always keep the largest A_k center in the case in which pruning implies I should remove all CBDs.

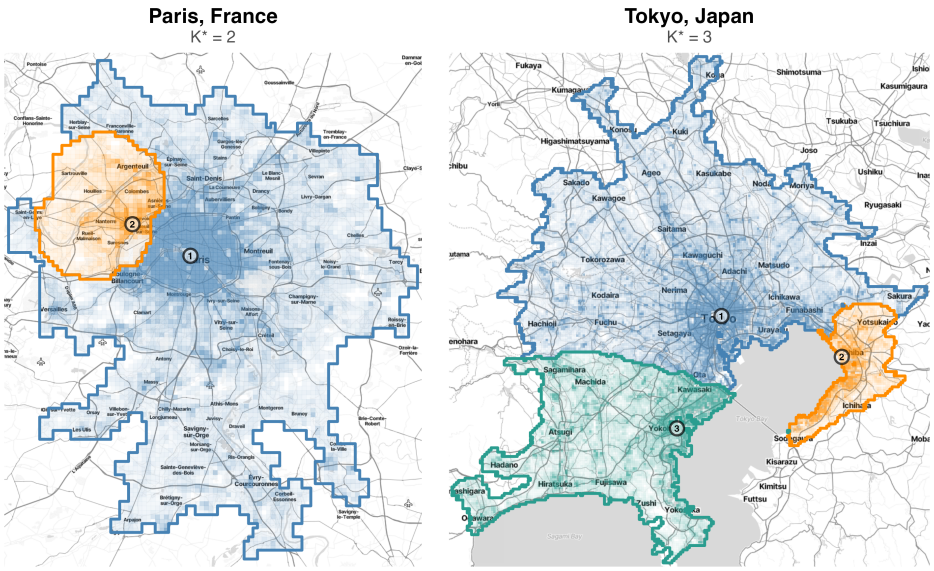


Figure 3: CBDs in Paris, France (left) and Tokyo, Japan (right). Pixels are aggregated to 500m and shaded by log built volume.

inversely proportional to the number of observations in its radial bin. This weighting scheme is designed to undo the mechanical rise in observations one obtains from the simple geometry of circles. Second, I exclude a small ring of observations around the CBD from estimation, where the size of the ring is generated to be proportional to the overall area of the city. I fix the scale of this 'donut' so that in a city the size of Shanghai (3,100km^2) I remove a 2km radius hole. I do this because in practice, CBDs are not pointless masses but have area that scales with the size of the city. Within CBDs, the building volume gradient is quite flat, which biases down θ_k. A summary of the CBD detection algorithm and θ_k estimation is in Figure 2.

2.3 Polycentricity around the world

The algorithm performs extremely well at finding well-known CBDs. For example, Figure 3 shows the CBDs detected in Paris, France ($K^* = 2$, one in central Paris and another surrounding La Défense), and Tokyo, Japan ($K^* = 3$, Chiyoda, Yokohama, and Chiba). By partitioning the UCDB agglomerations into potentially multiple cities, I expand the sample to 15,199 CBDs and their catchment areas, which I refer to as 'cities.'

In the left panel of Figure 4, I plot the average number of CBDs (K^*) in UCDB agglomerations within a country. Richer countries have more polycentric agglomerations: the average agglomeration in the US has

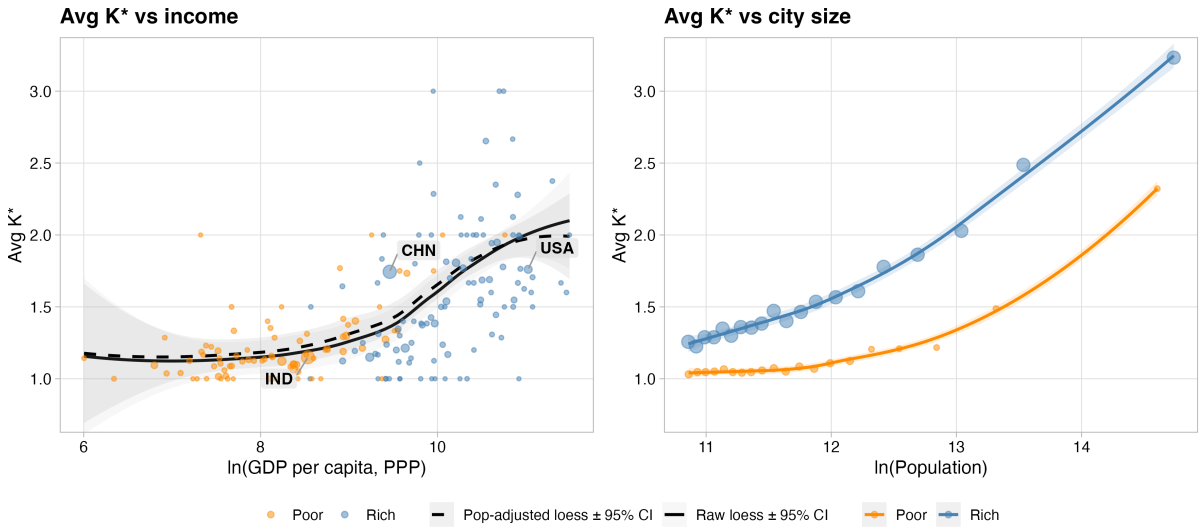


Figure 4: Average number of CBDs per agglomeration versus country-level log GDP per capita (left) and versus city size (right), split by the level of development.

about 1.75 central business districts, compared to 1.15 in India. This is not an artifact of simply the level of urbanization as this fact persists even when controlling for the size of the agglomeration. In the right panel, I plot the average K^* against log population of the agglomeration for both poor (World Bank 'Low' and 'Lower Middle' income countries) and rich countries. Cities in rich countries have more CBDs on average conditional on agglomeration size across the entire distribution of agglomeration sizes. Polycentricity is predominantly a large- and rich-city phenomenon.

2.4 City height, area, and density

Using the over 15,000 cities constructed by partitioning the UCDB agglomerations, I now examine how cities differ around the world in terms of physical aggregates like height and area, as well as average population density.

A city's average height is defined as the total built volume in the city divided by the amount of developable land, so this definition includes developed areas of a city with no height, like roads and parks, and not just the height of structures. Area is measured as the total developable area in km^2. Mechanically, I can decompose built volume per person into average height and average area per person.

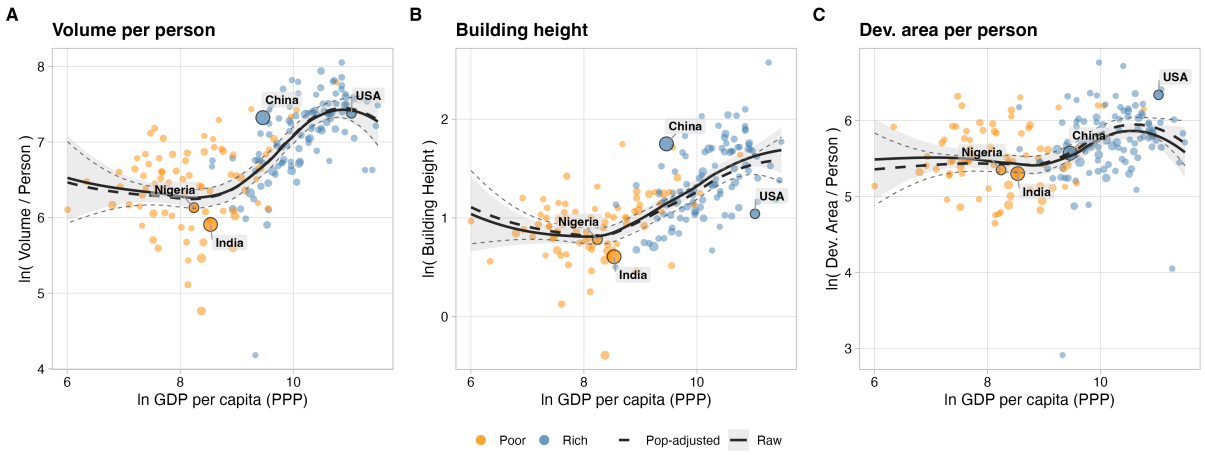


Figure 5: Average city characteristics versus a country's level of development. Left: built volume per person. Middle: average height per person. Right: developable area per person. $N = 15,199$ cities in 10,852 UCDB agglomerations.

In Figure 5, I loosely follow the approach of Jedwab et al. (2021) and ask how the characteristics of a city depend on its nation's level of development. Each scatterplot shows the simple average characteristic of cities within each nation versus the country's level of development as measured by (log) GDP/capita. Each plot shows both a lowess smoother of this relationship (solid line), and the relationship conditional on a city's size (dashed line).[15]

Cities in the developing world have substantially less built volume per person, on the order of one hundred log points or more.[16] This difference is primarily driven by height: rich world cities are considerably taller than their counterparts in poor nations, and are only modestly less dense.

In Appendix Table A1, I regress the difference between log city height and log developable area per person on country GDP/capita (with and without controlling for city size and primacy). When this coefficient is zero, it means the composition of built volume in terms of height or area is not systematically related to income. Instead I find that a 10% increase in GDP per capita shifts the composition of built volume per person towards vertical instead of horizontal construction by 1.1% ($p < 0.01$) indicating that cities in the developing world achieve their volume disproportionately with area rather than height: they build out, not

[15] That is, residualized against an estimated polynomial in city population and a dummy for whether a city is the largest in its country. I include this control because primate cities are often the focal site of government-led development, especially in nations with weak states (Ades and Glaeser, 1995).

[16] A one hundred log point difference here means that cities in the developing world have roughly 63% less built volume per capita than cities of a similar size in the rich world.

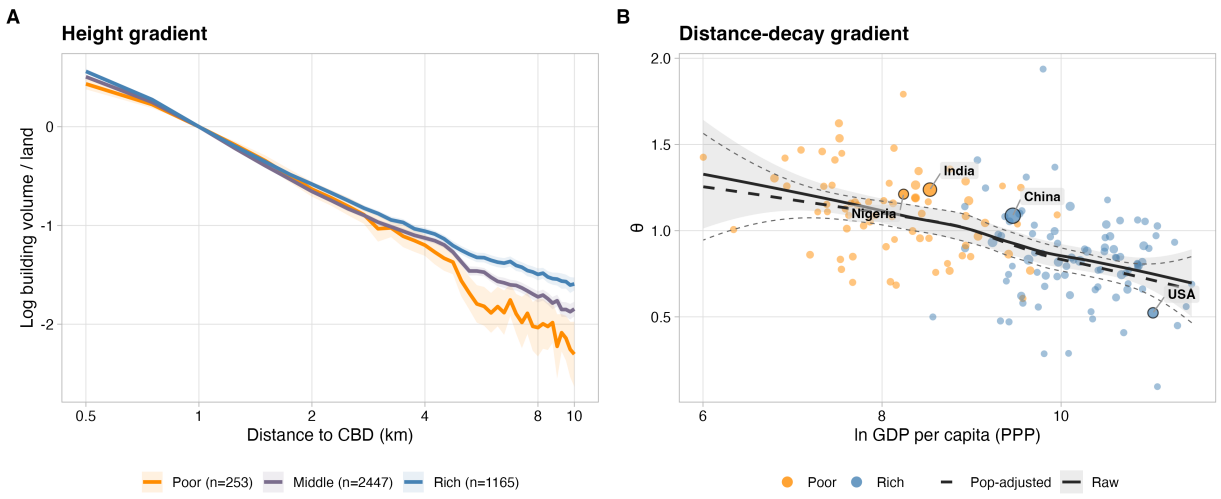


Figure 6: The left panel plots the average log building volume per unit land against distance bins of distance to the CBD, estimated by Poisson regression with city-cluster fixed effects and city-clustered standard errors. Distance bins are 250m wide and are presented on a log scale; the curve is normalized to zero at 1km. The sample excludes India and China, who have so many cities they dominate the curve. Panel B plots the distance-decay parameter θ_k (estimated at 100m resolution in an area-scaled donut around the CBD) against country-level log GDP per capita (PPP). Each point is a country average across non-primate cities; point size reflects the number of cities. The solid line is a loess smoother on raw country means; the dashed line adjusts for city population (quadratic polynomial residualization at the city level). Poor = Low + Lower Middle income; Rich = Upper Middle + High income (World Bank classification).

up. These results suggest that the technology to build cities varies around the world, and that there are either barriers to vertical development in the developing world, or horizontal barriers to development in the rich world.

2.5 Cities' internal structure

To better understand differences in cities' internal structure around the world, I examine estimates of θ_k for cities around the world.

In the left panel of Figure 6, I plot an estimate of cities' 'skyline slopes,' by which I mean the gradient of built volume per unit land with respect to log distance to the downtown. This provides a nonparametric estimate of how θ_k varies by the level of development. I average height in 250-meter concentric rings around each city's estimated CBD, and then I normalize height at a 1km distance from a city's estimated CBD (to remove level differences across cities made prominent in Figure 5). Averaging within distance bins removes

zeros, making a log-transform of the y-axis possible without issue. I then average across cities by the World Bank income group definition, omitting India and China, and plot 95% confidence bands for the estimate of the mean. The first thing to note is that log-linearity provides a fairly accurate description of cities' skyline slopes. Second, there appears to be a clear monotonic relationship between a country's level of development and the skyline slope of its cities. This pattern captures not only the height of structures in the urban core relative to the periphery, but also the thinning out of cities: relatively more urban land is undeveloped in the periphery. Building heights both decay more rapidly and cities thin out more quickly as you leave the downtown in developing cities compared to those in the rich world.

In the right panel of Figure 6, I plot estimates of the average θ_k across cities within a country against a country's level of development as measured by its log GDP per capita. I overlay a lowess smoother, both for the simple mean, and one in which I regression-adjust θ_k for differences in city size. The nonparametric relationship in the left panel persists with the parametric estimates of skyline slopes across the development spectrum: cities are flatter in the rich world.

Greater built volume density in the urban core relative to the periphery can reflect either a greater demand or supply for floorspace in the urban core, reflected by the fact that $\theta = \gamma \cdot \tau$: Demand for the urban core decreases with the quality of a city's transportation infrastructure. When commuting is costly, agents desire to live closer to the CBD, bidding up floorspace prices, incentivizing development, and ultimately resulting in a steeper skyline.

In summary, compared to cities in high-income nations, cities in the developing world accommodate demand for floorspace by building out, not up, yet pack in built mass relatively more in their urban cores, resulting in steeper skylines.

3 A quantitative model of an urban system

I introduce a general equilibrium model rich enough to show how differences in urban form affect aggregate outcomes and simple enough to be estimated using the limited geospatial data available for all cities in the world. My goal is to endogenize both average density, height, and the area of cities, as well as the spatial

distribution of floorspace within cities, and provide a framework to understand the aggregate impact of these city-level differences that can be quantified with these data. In the model, each country is a closed economy consisting of monocentric cities and an agricultural sector.[17] Agents choose whether to live in a city, in which city to live, and where in that city to live. A continuum of competitive developers construct floorspace and urban land. Cities produce urban varieties traded across cities, while the agricultural sector produces the numeraire good.

In what follows, I layout the model environment, preferences, technology, and equilibrium definition, followed by some discussion on the assumptions made. I then characterize the equilibrium.

Environment A nation n is composed of I_n cities indexed by i, with the index $i = a$ denoting the agricultural sector. There is a mass L_n of ex-ante identical households that choose whether to work in the agricultural sector or urban sector, and conditional on working in the urban sector, decide in which city to live and their residency location (x, ϕ) given in polar coordinates, where x denotes distance to city's central business district (CBD). As the model is of a single nation in autarky, I omit n subscripts in this section.

Cities are circular and have an endogenous radius X_i. I restrict my attention to symmetric allocations – i.e., at each x, allocations are identical over ϕ – so it is sufficient to consider the location choice over x. All urban production takes place at the CBD, to which agents commute for work.

Households earn a wage w_i, which finances their consumption of urban varieties, the agricultural good, and floorspace. Competitive developers construct floorspace $H_i(x)$ on each unit of land, after paying a fixed cost to incorporate marginal land into the city. The urban radius X_i is pinned down when land rents are equal to the fixed cost of urban land development. Urban development is absent in the agricultural sector, as is the need to commute. Instead, in the agricultural sector, agents' demand for floorspace is met by directly consuming land; they live on the farm. This assumption means that in the agricultural sector, agents earn

[17] I model nations in autarky to abstract from the role trade plays in the urbanization process. In my model, international trade would mediate how goods prices move in counterfactuals, depending on the extent to which goods markets clear nationally or internationally. Many developing countries are agricultural exporters, so price fluctuations driven by rural-to-urban migration would be dampened by including international trade. This would amplify the welfare effects of rural-to-urban migration I study in Section 5. Some work has theorized on the effect of trade openness on urbanization. Trade may impede urbanization by encouraging specialization in the production of agricultural goods; indeed, the world's least urbanized countries tend to be agricultural exporters. Trade may also incentivize clustering at international gates (Hanson, 1998; Coşar and Fajgelbaum, 2016) raising the size of, e.g., port cities (Ducruet et al., 2024). Nagy (2022) finds that cities close to borders shrank after those borders moved in Hungary. I capture international market access through a city-specific productivity term. As I am not interested in counterfactuals that change trade costs, or spillovers from urbanization through the trade network, it is appropriate to hold these productivities fixed. In my view, developing quantitative models of international trade and urbanization is a compelling path for future research.

their average product, as they own the land which they till. This is the assumption maintained in Lewis (1954) and the formal models that it inspired; see Gollin (2014) for a review of the Lewis model and its applications.

Household preferences Households, indexed by ν, have Cobb-Douglas preferences over urban varieties, agriculture, and floorspace. Alongside goods demand, they choose a city i or the agricultural sector in which to live and work, and, if choosing to live in a city, they choose a distance x from the city's CBD at which to live. Agents' location choice depends on a realization of an idiosyncratic preference shock, which is iid across cities and the agricultural sector, but households do not receive preference shocks for locations within cities. They inelastically supply labor and earn a wage w_i if they live in city i. They solve,

$$\max_{\{c_j\},h,i,\ x\in(0,X_i)} A_i(x)\left(\frac{C}{\alpha}\right)^{\alpha}\left(\frac{\psi^H h}{\beta}\right)^{\beta}\left(\frac{c_a}{1-\alpha-\beta}\right)^{1-\alpha-\beta}\epsilon_i^{\nu}, \quad C=\left(\sum_{j=1}^{I} c_j^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \tag{4}$$

subject to the budget constraint $\sum_{j=1}^{N} p_{ji}c_j + c_a + q_i(x)h \leq w_i$, where p_{ji} is the price of city j's variety in city i. The price of floorspace, $q_i(x)$, depends on the city and location x within a city in which an agent resides. The term ψ^H reflects a quality-adjustment factor reflecting that while agents pay for units floorspace h, they enjoy it depending on its quality, which I assume to be constant across cities within a country. I include this quality adjustment term so that physical floorspace (observed in the data) does not capture quality-adjusted floorspace, which may vary across countries.

An agent's demand for floorspace in city i at x is denoted, $h_i(x)$, while in the agricultural sector, agents' demand is denoted by h_a. c_{ij} denotes demand for urban variety i by an agent that chooses to locate in j. An agent's indirect utility for location x in city i is $v_i(x)$, which will be equalized across all x within each city i to a common level, v_i. Indirect utility for agents living in the agricultural sector is denoted v_a.

The term $A_i(x)$ is a city-site specific amenity. The term ϵ_i^{ν} is an idiosyncratic preference shock drawn iid across locations from a Fréchet distribution with location parameter 1 and shape parameter ε.

My treatment of preferences and budget constraints departs from standard assumptions in the system-of-cities literature since Henderson (1974). Much of this literature assumes each household consumes one unit of land; I allow households to flexibly adjust the amount of floorspace they consume and fix an income elasticity of floorspace consumption to 1. Estimates of the income elasticity of housing demand range from

unit elastic to mildly inelastic. In the U.S., both Davis and Ortalo-Magné (2011) and Comin et al. (2021) estimate a number near 1, as implied by Cobb-Douglas, while Albouy et al. (2016) and Finlay and Williams (2022) estimate an elasticity closer to 0.7. In China, Murray and Sun (2017) estimate the same elasticity. Malpezzi (1999) surveys estimates of housing demand price and income elasticities in developing economies from earlier literature that broadly accord with these patterns.

Commuting costs Cities are exogenously characterized by their amenities supply function,

$$A_i(x) = \tilde{A}_i \cdot (x)^{-\tau_i},$$

where $\tilde{A}_i$ reflects a common, exogenous citywide amenity A_i, multiplied by city-specific constants; see Appendix B.[18] The component of amenities given by $x^{-\tau_i}$ reflects commuting costs. I have assumed that commuting costs are paid in utils, not dollars, to capture decreases in other amenities beyond the opportunity cost of travel. Assuming wages deteriorate at rate τ_i does not meaningfully impact the model's predictions. Cities with a higher τ_i have a larger distance elasticity of commuting costs, reflecting worse urban transportation technology. This formulation abstracts from consumption differences across space driven by trade costs within the city.

Urban development technology In each city, there is an infinite pool of identical potential developers. Developers choose whether to enter, and upon entering, operate a technology at a single site to build floorspace vertically after paying a fixed cost to incorporate marginal land into the city. The technology available to builders for building floorspace H_i at distance x is,

$$H_i(x) = \frac{\widetilde{Z_i^H}}{\psi^H} c_a^{\frac{\gamma_i}{1+\gamma_i}} T_i(x, \phi)^{\frac{1}{1+\gamma_i}}. \tag{5}$$

At each location, available land is fixed $T_i(x, \phi) = 1$. The term $\frac{\widetilde{Z_i^H}}{\psi^H}$ captures fundamental productivity in building floorspace, Z_i^H, and a quality-adjustment term ψ^H, multiplied by city-specific constants that ensure the intercept of the floorspace supply curve does not depend on γ_i, as clarified in Appendix B. Fundamental productivity in floorspace construction may differ across cities and captures one aspect of the level of urban costs in a city. The quality adjustment term ψ^H is the same term that appears in the utility function and is

[18] $\tilde{A}_i$ reflects A_i multiplied by a city-specific adjustment factor which ensures that the city specific elasticities (γ_i, τ_i) only affect how utility scales with population, and not the effective *level* of utility in each city.

assumed to be constant across cities within a nation.

I assume a Cobb-Douglas production function for building floorspace vertically, which results in a constant-elasticity supply function where the supply elasticity is γ_i. There is good evidence the production function for floorspace is Cobb-Douglas. Epple et al. (2010), Ahlfeldt and McMillen (2014), and Combes et al. (2021) all estimate the housing production technology in a variety of contexts and find that a constant expenditure share on land and a elasticity of substitution between land and materials near one fits the data well. Heterogeneity in γ_i across cities captures differences in geophysical and regulatory barriers to vertical development.[19]

To incorporate marginal land into the city, developers must pay a fixed cost $F_i(x)$ that is rising in x,

$$F_i(x) = \frac{\widetilde{Z_i^X}}{\psi^X}(x)^{2/\rho_i}$$

where $\frac{\widetilde{Z_i^X}}{\psi^X}$ is a scaling of fundamental productivity in incorporating land into the city, Z_i^X and a quality adjustment term, ψ^X, constant across cities in a nation, multiplied by city-specific scaling constants; see Appendix B. Developers are competitive and take prices as given. A developer constructing floorspace at x solves,

$$\max_{c^a} q_i(x)H_i(x) - c^a - F_i(x), \tag{6}$$

subject to the production technology (5). Developers only enter and construct floorspace so long as land rents, $r_i(x) \geq F_i(x)$, where, $r_i(x) = \max_{c^a} q_i(x)H_i(x) - c^a$.

The assumption that developers can build the city out, subject to increasing marginal costs is nonstandard. In the quantitative spatial literature, land supply is typically assumed to be fixed (i.e., $\rho_i = 0$). In classical urban models, the canonical assumption is that land is available everywhere at some constant reservation price (i.e., $\rho_i = \infty$). I instead treat the *land supply elasticity* ρ_i, which is the elasticity of urban area supplied with respect to land rent at the periphery, as a parameter to be estimated.[20]

[19] Gomtsyan (2024) also shows that financial constraints can hold back the construction of tall buildings in developing cities.

[20] In quantitative models of cities like those using the framework of Ahlfeldt et al. (2015), this assumption manifests through assuming the (potential) number of sites workers and firms can occupy is held fixed. My assumption on the production of urban land is a 'reduced form' representation of the urban developer's problem. One microfoundation is that the easiest-to-develop locations are built first, and so as a city expands, developers are incentivized to move up the land supply curve and build atop increasingly difficult to develop land. In Los Angeles, for example, this means that an expansion of the urban extent requires building into the Hollywood Hills, or for a city like Singapore to reclaim land from the sea. It may also reflect regulatory constraints, in addition to geophysical ones. For example, both London and Seoul's development is inhibited by a greenbelt that surrounds the city.

Total development profits are,

$$\Pi_i = 2\pi \int_0^{X_i} x\big(r_i(x) - F_i(x)\big)dx.$$

Goods production Urban varieties are produced with a technology that exhibits external scale economies in density,

$$y_i = Z_i^y L_i, \quad Z_i^y = \bar{Z}_i^y \left(\frac{L_i}{\pi X_i^2}\right)^{\zeta}.$$

$\bar{Z}_i^y$ is the exogenous component of goods productivity in city i. Letting $\zeta > 0$, the endogenous component of productivity represents myriad agglomeration economies present in cities like knowledge spillovers and labor market pooling, which here are modeled as technological externalities. Urban producers are competitive in goods and labor markets, and choose L_i to maximize profits. Urban varieties are traded across locations subject to iceberg trade costs $\delta_{ij} \geq 1$, where I assume $\delta_{ii} = 1$.

My assumption on the urban goods technology varies slightly from standard representations in the spatial literature. While I assume that external economies depend on density, which is a standard assumption, here *density*, and not just urban population, is endogenous.[21] The externality reflects two agents failing to internalize their effect on urban aggregates. First, urban goods firms under-hire, as they fail to internalize the fact that more workers will increase city productivity (Z_i^y) by engaging in productive social interactions and facilitating knowledge spillovers. Consequently, workers' wages are too low. Second, atomistic developers fail to account for the fact that by spreading people out over a larger area, they reduce the number of productive social interactions that occur, lowering their rents in equilibrium; land rents at the periphery are too high relative to a social optimum. This feature of the model captures how differences in the urban technology across space affect the benefit of scaling cities in size, even when ζ is constant across cities, as the elasticity of density with respect to population, holding wages fixed is $1/(1+\rho_i)$. This feature of the model does not prove to be quantitatively important in Section 5.

[21] Duranton and Puga (2020) argue that *experienced* (i.e., population-weighted), not average density is the appropriate measure for understanding urban scale economies. My model gives a closed-form expression for this quantity as well,

$$\underbrace{2\pi \int_0^{X_i} x \frac{L_i(x)}{L_i} L_i(x) dx}_{\text{experienced density}} = \left(1 - \tau_i \frac{1+\gamma_i}{\beta}\right) X_i^{-\tau_i \frac{1+\gamma_i}{\beta}} \cdot \left(\frac{L_i}{\pi X_i^2}\right).$$

Thus, in my framework, experienced density is a scaling of average density that depends on the area of the city and parameters that govern the urban technology. This leads to a direct effect of transportation technology (τ_i) and building technology (γ_i) parameters on productivity, in addition to endogenous effects through how it affects the equilibrium average density in the city.

In the rural sector, agricultural goods are produced with a constant returns technology in land and labor,

$$y_a = Z_a(L_a)^{1-\mu}(T_a^y)^{\mu},$$

and are freely traded.[22] Agricultural producers choose L_a and T_a^y to maximize profits. I assume that agricultural workers own the land on which they work, and so their compensation reflects wage income and rebated land rents.

Equilibrium definition Given urban and rural fundamentals, $\{\bar{A}_i, Z_i^H, Z_i^X, \bar{Z}_i^y, \psi^H, \psi^X\}$, city-specific technological parameters $\{\tau_i, \gamma_i, \rho_i\}$, production parameters $\{\zeta, \mu\}$, preference parameters $\{\alpha, \beta, \sigma\}$, and trade costs $\{\delta_{ij}\}$, an equilibrium is a population distribution across locations $\{L_i\}$, across sites in cities $\{L_i(x)\}$, city radii $\{X_i\}$, alongside floorspace prices $\{q_i(x)\}$, goods prices $\{p_i\}$, wages $\{w_i\}$, and a common level of indirect utility within each city $\{v_i\}$, such that,

1. households, taking wages and prices as given, optimally choose i, x (if choosing a city), alongside floorspace and goods demands, solving (4);
2. each developer, taking floorspace prices as given, solves (6);
3. developers enter at each x until land rents are equal to the fixed cost of land incorporation, so that at X_i, $r_i(X_i) = F_i(X_i)$;
4. the floorspace market clears at each (x, ϕ) in every city,

$$H_i(x) = h_i(x)L_i(x);$$

5. all urban households are housed somewhere,

$$2\pi \int_0^{X_i} xL_i(x)dx = L_i;$$

6. a spatial equilibrium holds within each city, so that (indirect) utility is equalized across all sites in the

[22] This assumption is motivated theoretically by the fact that agricultural production can occur anywhere on land that is not urban, and empirically by the fact that in my data, I do not observe the spatial distribution of agricultural activity and cannot recover trade costs.

city,

$$v_i(x) = v_i \quad \forall x \in (0, X_i], \quad \forall i = 1, ..., N$$

7. urban production firms, taking wages and prices as given, optimally choose labor demand;

8. agricultural producers choose L_a and T_a^y to maximize profits;

9. the labor market across cities and the agricultural sector clears,

$$\sum_{i=a,1,...,I} L_i = L;$$

10. land rents in agriculture are rebated back to agricultural households;

11. the land market clears in the agricultural sector, so that,

$$T_a = T_a^y + h_a L_a;$$

12. and the goods market clears all urban varieties, so that,

$$y_i = \sum_{j=1}^{I} \delta_{ij} c_{ij} L_j + \delta_{ia} c_{ia} L_a,$$

and development profits are consumed by absentee landlords on the numeraire good so that when the agricultural goods market clears,

$$Y_a = \sum_{j=1}^{I} \left(2\pi \int_0^{X_j} x c_a(x) dx + \Pi_j + c_{aj} L_j \right) + c_{aa} L_a,$$

where $c_a(x)$ reflects the demand from the developers constructing floorspace at distance horizon x.

Discussion My model rationalizes equilibrium gaps between agricultural and urban wages through several mechanisms. First, the urban cost of living may be higher than it is in the agricultural sector, so that rural-urban real wage gaps may be smaller than differences in nominal wages. Second, agricultural amenities may be higher, which keeps workers in the agricultural sector. Agricultural amenities capture a variety of forces that keep workers in the agricultural sector. For example, they may capture informal insurance through social

networks present in the agricultural sector (Munshi and Rosenzweig, 2016). Additionally, these amenities may reflect a wedge driven by migration costs.[23] Finally, the parameter ε, the migration elasticity, captures how responsive migration flows are to changes in amenity-adjusted real wages across space. If the migration elasticity is not infinite, the model can rationalize amenity-adjusted real wage gaps between sectors as existing in equilibrium, as it assumes there is a mass of inframarginal agents enjoying their idiosyncratic component of utility in each city and the agricultural sector. However, I do not consider urban unemployment as an adjustment mechanism, as in Harris and Todaro (1970), as I do not have global data on urban unemployment.

3.1 Model solution

City structure I first outline a city *partial* equilibrium in which p_i, w_i, and L_i are fixed (i.e., equilibrium conditions 1-6) and then outline the model's general equilibrium (conditions 7-12). All derivations are detailed in Appendix B.

Combining households' demand for floorspace (equilibrium condition 1) with the definition of a city's spatial equilibrium (condition 6) determines the gradient of floorspace prices,

$$q_i(x) = \left(\frac{A_i(x) w_i P_i^{-\alpha}}{v_i} \right)^{1/\beta} \quad \forall x \in (0, X_i). \tag{7}$$

Solving the developer's problem (6), floorspace supply is isoelastic in its price, namely,

$$H_i(x) = \frac{Z_i^H}{\psi^H} q_i(x)^{\gamma_i}. \tag{8}$$

Using (7) and (8), the floorspace market clearing (condition 4), and labor market clearing (equilibrium condition 5), it can be shown that the spatial distribution of population within cities is given by,

$$L_i(x) = \left(1 - \frac{\tau_i}{2} \frac{1+\gamma_i}{\beta} \right) \left(\frac{x}{X_i} \right)^{-\tau_i \frac{1+\gamma_i}{\beta}} \frac{L_i}{\pi X_i^2}. \tag{9}$$

Population density falls in space at rate $-\tau_i \frac{1+\gamma_i}{\beta}$. Cities with more elastic floorspace supply ($\gamma_i \uparrow$) or large

[23] When migration costs are not bilateral, but instead are depend on an origin and destination component, they are not separable from amenities, as in Desmet et al. (2018).

commuting costs ($\tau_i \uparrow$) are denser in the core relative to the periphery. Absent commuting costs, population density would be uniform over space.

Finally, to pin down X_i, and using the equilibrium value of land rents ($r_i(x)$), by equilibrium condition (3), developers enter until land rent at the urban periphery equals the marginal cost of development, pinning down the urban radius X_i,

$$\frac{\beta w_i L_i(X_i)}{1+\gamma_i} = \frac{\widetilde{\psi^X}}{Z_i^X}(X_i)^{2/\rho_i}. \tag{10}$$

Equations (7)-(10), give solutions to $q_i(x), L_i(x), H_i(x), v_i$ and X_i conditional on p_i, w_i and L_i. Prices, and the population distribution across cities and the agricultural sector are determined in general equilibrium. First, profit maximization from urban goods producers (condition 7) ensures that urban wages reflect marginal revenue products, $w_i = p_i Z_i^y$.

I assume that land rents are rebated back to households in the rural sector; i.e., they own the land on which they work, so that their compensation reflects both the average product of labor in agriculture and their share of land rents,

$$w_a = (1-\mu)\frac{Y_a}{L_a} + \frac{q_a T_a}{L_a}.$$

In the agricultural sector, land is consumed directly instead of floorspace. Production firms demand T_a^y and agricultural households demand $h_a L_a$, which in equilibrium must sum to the total amount of land available, T_a. When T_a^y is chosen optimally so that the marginal product of land in production equals its price, then in equilibrium, land prices must reflect a scaling of the marginal product of land,

$$q_a = \left(\frac{\beta(1-\mu)+\mu}{1-\beta}\right)\frac{Y_a}{L_a},$$

reflecting its dual uses in production and consumption.

Finally, in equilibrium,

$$v_i = A_i (Z_i^H)^{\frac{\beta}{1+\gamma_i}} (Z_i^X)^{\frac{\beta}{1+\gamma_i}-\frac{\tau_i}{2}} (\psi^H)^{\beta\frac{\gamma_i}{1+\gamma_i}} \frac{w_i}{P_i^\alpha} (\beta w_i L_i)^{-\kappa_i}$$

$$v_a = \hat{A}_a \frac{w_a}{P_a^\alpha} (w_a L_a)^{-\beta},$$

where term $\hat{A}_a$ combines agricultural amenities with other model constants; see Appendix B. Furthermore,

$$\kappa_i = \frac{1}{1+\rho_i}\frac{\beta}{1+\gamma_i} + \frac{\rho_i}{1+\rho_i}\frac{\tau_i}{2}, \tag{11}$$

is the *urban cost elasticity*, and it captures the costs of urban scale in city i. It measures how consumption-equivalent utility falls in L_i, holding wages and traded goods prices fixed.

To determine L_i, I use the assumption of Fréchet distributed idiosyncratic preference shocks. The population distribution across cities and the rural sector is given by,

$$L_i = \frac{(v_i)^{\varepsilon}}{\sum_{j=1}^{I}(v_j)^{\varepsilon} + (v_a)^{\varepsilon}} L_n.$$

As prices reflect factory gate prices gross transportation costs, $p_{ij} = \delta_{ij} p_i$, from goods market clearing for urban goods, we have,

$$w_i L_i = \sum_{j=1}^{I} \left(\frac{\delta_{ji} p_i}{P_j}\right)^{1-\sigma} \alpha w_j L_j + \left(\frac{\delta_{ia} p_i}{P_a}\right)^{1-\sigma} \alpha w_a L_a. \tag{12}$$

Since the share of urban income spent on floorspace β is used either to purchase construction materials from the agricultural sector, or transformed into profits used to buy agricultural goods, goods market clearing for the agricultural sector can be written as,

$$(1-\alpha)\left(\sum_{i=1}^{I} w_i L_i + w_a L_a\right) = Y_a.$$

Provided (12) holds for each city i, then by Walras' law the above equality holds.

Equilibrium existence and uniqueness Given the presence of increasing returns, there is a possibility of multiple equilibria. However, the congestion forces in the model are sufficient enough to rule this out.[24] A

[24] This can be shown numerically using Theorem 1 of Allen et al. (forthcoming). The theorem of that paper is a sufficient condition for existence and uniqueness. It states that if a square matrix whose ijth element bounds above the absolute value of the elasticity of one type of endogenous variable with respect to another (here, wages and population) has a spectral radius less than or equal to one, the equilibrium exists and is unique. Numerically, this amounts to showing that agglomeration forces (ζ) must be weak relative to congestion forces (κ_i, ε). While the model does not yield a simple parameter restriction, the condition can be numerically checked. Appendix B.1 details this argument and provides a numerical check that the calibrated model has a unique equilibrium.

necessary condition for existence is,

$$\frac{\beta}{1+\gamma_i} > \frac{\tau_i}{2} \quad \forall i.$$

This condition restricts the effects of land on city-level outcomes. Consider a city with an exogenous amount of land, $\rho_i = 0$, and consider a change dX_i. Increased land availability has two effects. First, by spreading out agents over a larger area, there is less housing market congestion, increasing city welfare at rate $2\frac{\beta}{1+\gamma_i}$.[25] However, by increasing the radius, an agent residing at the periphery is dX_i/X_i percent farther from the core and pays $\tau_i\%$ more in commuting costs. This condition says that the net effect of land must have a positive effect on city welfare. Consider a situation in which this inequality did not hold. This could not be an equilibrium: agents would not want to spread out over the entire area provided, since the gain from reduced housing market congestion would not outweigh the loss from longer commutes. Instead, agents would deviate to a smaller area: the city would collapse over an area of measure zero at $x = 0$. Using the estimates from Section 4, I empirically find this condition holds for 96% of cities in the sample.[26]

The urban cost elasticity The key measure that governs how well a city can absorb total city income is the urban cost elasticity, κ_i, defined in equation (11). It is the elasticity of city consumption-equivalent welfare with respect to city population, holding wages and goods prices fixed.[27] That is, it answers the question, if a city grew 1% in population, what percent increase in consumption would be required to keep residents as well off as they were before, holding wages and goods prices fixed? It is important to hold wages and goods prices fixed, because changes in these prices may reflect the benefits of urbanization, through, for example, agglomeration effects, or spillovers through the trade network that engender increased market access. The urban cost elasticity contains all components of the urban technology: how well cities can build up, build out, and how costly residents find it to spread out. To gain intuition on how it operates, consider two edge cases,

$$\rho_i = 0 \implies \kappa_i = \frac{\beta}{1+\gamma_i}, \quad \text{and} \quad \rho_i = \infty \implies \kappa_i = \frac{\tau_i}{2}.$$

When land supply is completely inelastic ($\rho_i = 0$) as a city grows in population, land is increasingly scarce. The ability for a city to accommodate people is determined entirely by its ability to build 'up' (γ_i) and so all

[25] The factor 2 arises due to basic geometry: area is proportional to the square of radius.
[26] For cities where the condition does not hold, I shrink τ_i so that condition holds, only in the quantified model.
[27] This is the definition in Combes et al. (2019), which derives this object using a different model.

congestion forces operate through the floorspace market. When land supply is completely elastic ($\rho_i = \infty$), any upwards pressure in land prices from a population increase incentivizes developers to expand the city horizontally so as to completely offset congestion through the housing market. The consequence of a larger city is that commutes are on average longer, lowering resident welfare.

4 Estimating the urban technology

I seek to estimate the elements of the urban technology that determine city structure and the urban cost elasticity. These are the commuting cost elasticity, τ_i, the floorspace supply elasticity, γ_i, and the land supply elasticity ρ_i. To do so, I rely on three moments in the data that, through the structure of the model, identify these parameters.

Using conditions (8), (7), the floorspace market clearing (condition 4), and plugging in for the population distribution (9), yields

$$H_i(x) = \frac{Z_i^H}{\psi_n^H} \left(\frac{\beta w_i L_i}{\pi X_i^2} \left(1 - \frac{\tau_i}{2} \frac{1+\gamma_i}{\beta} \right) \left(\frac{x}{X_i} \right)^{-\tau_i \frac{1+\gamma_i}{\beta}} \right)^{\frac{\gamma_i}{1+\gamma_i}}, \tag{13}$$

which informs how the floorspace per unit of land gradient, $H_i(x)$, depends on a city's commuting costs, τ_i, and its floorspace supply elasticity, γ_i. Cities with higher commuting costs or more elastic floorspace supply have relatively denser urban cores than peripheries. When commuting costs are high, demand for land in the urban core is high as agents seek to avoid long commutes.

Integrating over all sites in city i in nation n gives,

$$\underbrace{2\pi \int_0^{X_i} x H_i(x) dx}_{\equiv H_i} = 2\pi \frac{Z_i^H}{\psi_n^H} (\beta w_i L_i)^{\frac{\gamma_i}{1+\gamma_i}} (\pi X_i^2)^{\frac{1}{1+\gamma_i}} \tilde{\tau}_i, \tag{14}$$

where $\tilde{\tau}_i$ is a nonlinear function of τ_i and γ_i. This equation describes how γ_i controls how log total built volume (H_i) of a city scales with its log income and log area.

Third, using developer optimality for X_i, equation (10), and the equilibrium value of floorspace profits at the

periphery reveals that,

$$\pi X_i^2 = \pi \frac{Z_i^X}{\psi_n^X} \left(\frac{\beta}{1+\gamma_i} w_i L_i \left(1 - \frac{\tau_i}{2} \frac{1+\gamma_i}{\beta} \right) \right)^{\frac{\rho_i}{1+\rho_i}}. \tag{15}$$

This equation describes how urban land scales with urban income at a rate that depends on ρ_i.

I use these relationships as a basis to estimate τ_i, γ_i, and ρ_i, which requires data on cities' built volume, area, and income. First, I estimate $\tau_i \gamma_i / \beta$ for each city using data on the internal structure of cities. Second, I use the cross-sectional relationship between city income, built volume, and footprint to estimate γ_i with instrumental variables. To estimate ρ_i, I exploit the time series of urban growth in income and area.

4.1 Data

The following uses the urban agglomerations in the Global Human Settlement Layer's (GHSL) Urban Centre Database (UCDB) (2019 version, Florczyk et al., 2019), and the GHSL's definition of built volume, and Google Maps' queried downtowns, opposed to the data constructed in Section 2. The reason is **this section is undergoing substantial reconstruction** and the analysis will, in future drafts, be completely re-done using the larger sample of cities and measurement of skyline slopes.

After dropping observations flagged for potential poor quality, I am left with 10,303 cities of over 50,000 persons. I associate with each city an auxiliary data product from the GHSL: built volume in 2015, which uses a variety of satellite sources and digital elevation models to construct building height and surface area at the 100m$\times$100m resolution within each urban agglomeration, and leaves me with over 300 million observations of built volume within cities globally.

The UCDB data product contains time series data on urban area (in km^2), urban population (from Gridded Population of the World, which provides the most spatially granular census data available) and GDP estimates at the city level from Kummu et al. (2018) for years 1990, 2000, and 2015. The GDP data is simply sub-national GDP data that is extrapolated to the city level by population-weighting according to the size of the city, and is likely measured with considerable error. I instead construct an alternative measure of city GDP

in 2015 using the Visible Infrared Imaging Radiometer Suite (VIIRS) nighttime luminosity measures.[28] To estimate city GDP, I use the 2015 cross-sectional relationship between log GDP and log nightlights, conditional on country fixed effects in Europe and North America, where the sub-national data is sufficiently granular, and use this relationship as a basis to predict urban GDP for cities outside of that sample. This routine is similar in spirit to using nightlights to measure economic growth when national accounts data are poorly measured (Henderson et al., 2012) and is approximately the same routine used in Ducruet et al. (2024) to predict city GDP across port cities. I aggregate my GDP measures so that they sum to the national income (in 2015 USD) net of agriculture and housing value-added. I also measure the DMSP-OLS nightlights for each city, as its time series goes back to 1992, which I associate with year 1990 in the data.[29]

Using Google Earth Engine, I compute for each 100m$\times$100m pixel the share of land within it that is developable. I classify land as developable following Saiz (2010): land is developable if it is not liquid water, ice, herbaceous wetland, and has a slope of less than 15%. For landcover classification, I use the European Space Agency WorldCover data (Zanaga et al., 2022), which has a 10m resolution. I use the NASA Digital Elevation Model, which is available at the 30m resolution (NASA JPL, 2020), to recover slope and elevation measures. For each city, I also recover the average clay, water, and sand content of city soil alongside soil density at 1m depth using the OpenLandMap. I detail all data sources and my extrapolation procedure for city GDP in Appendix C.

4.2 Built volume gradients

Taking logs of (13) gives a relationship between built volume at (x, ϕ) and its distance to the CBD,

$$\frac{d \log H_i(x, \phi)}{d \log x} = -\frac{\tau_i \gamma_i}{\beta}. \tag{16}$$

For data on $H_i(x, \phi)$, I use built volume at the 100m$\times$100m resolution for all cities in the sample, divided by the amount of developable land in that cell.[30] For distance to the CBD, x, I use the straight-line distance

[28] A common alternative nightlights product is the Defense Meteorological Satellite Program (DMSP) Operational Linescan System (OLS). I prefer the VIIRS for GDP extrapolation, as the VIIRS are about half as noisy as the DMSP-OLS nightlights time series, unmarred by top-coding and blurring issues, and perform considerably better at predicting city GDP (Gibson et al., 2021).

[29] However its time series ends in 2012.

[30] This is the right notion of height to use, as it captures floorspace per unit land, consistent with the theory. Rich cities have skyscrapers in their downtowns, giving them 'pyramid' shaped skylines as in Lall et al. (2021b) and Lall et al. (2021a), who use

between a cell's centroid and the coordinates associated with Google Map's definition of the city center.[31] For cities that do not appear in Google Maps, I use the centroid of the polygon that defines the urban extent.[32]

For each city, I estimate the sample analog of (16),

$$H_i(x,\phi) = \exp\left(-\frac{\tau_i\gamma_i}{\beta}\log x + \xi_{i,\phi}\right) t_i(x,\phi), \tag{17}$$

where $t_i(x,\phi)$ is the error term that captures measurement error in $H_i(x,\phi)$ and I assume that $\mathbb{E}[t_i(x,\phi) \mid \log x, \xi_{i,\phi}] = 1$. $\xi_{i,\phi}$ are city-by-polar-angle fixed effects. As the model traces out $-\frac{\tau_i\gamma_i}{\beta}$ by examining how $H_i(x,\phi)$ changes in x, holding ϕ fixed, I hold ϕ fixed by examining density gradients within bins of ϕ within the city. Moreover these fixed effects capture the overall level of $H_i(x,\phi)$ within the city. For a circular city, the number observations mechanically rises as you leave the downtown, which could bias estimates towards the built volume slope farther away, if the slope is spatially heterogeneous. I enforce that observations at any distance equally contribute to the measure of the slope by reweighting observations. For each city, I create 100 bins of x and enforce that the weight of an observation is inversely proportional to the number of observations in a bin. I restrict observations to be within 10km of the CBD. Restricting observations to 10km drops overlap with the majority of 'satellite cities.' The built volume data is mapped in the left panel of Figure 7 for Shanghai. Built volume decays rapidly from the CBD.

I use a pseudo-Poisson maximum likelihood estimator with fixed effects (Wooldridge, 2010; Correia et al., 2020) to estimate equation (17). The reason is that there are zeroes in $H_i(x,\phi)$ due to the presence of, e.g., parks within cities. Dropping these would be inappropriate – consistent with the theory, empty lots are more likely to appear in the periphery of cities than at the core as the incentives to develop there are smaller. For example, Hongqiao International Aiport is a large grey strip south-southwest of the CBD in Figure 7. Dropping these observations would thus give a biased built volume gradient.

pixel heights from the World Settlement Footprint. However, this does not necessarily reflect more built mass in their downtowns, as there remains developable land that has not been built up.

[31] It is common in the urban literature to measure distances within cities using drivetimes or distances on the road network. That is inappropriate in this exercise, precisely because deviations from the straight-line distance are what I seek to capture in τ_i. That is, the worse a city's network can replicate straight-line distances, the worse transportation technology it has. The notion of transportation technology here is broad, and not just limited to road speed or the presence of public transportation infrastructure. For example, a city like Seattle has many bodies of water for which crossing is costly, disincentivizing sprawl.

[32] 27% of the cities in the data do not appear when queried by Google Maps, either because their name is missing, given in the local writing system, or simply do not appear. Such cities are about 40% smaller in terms of population than the average city.

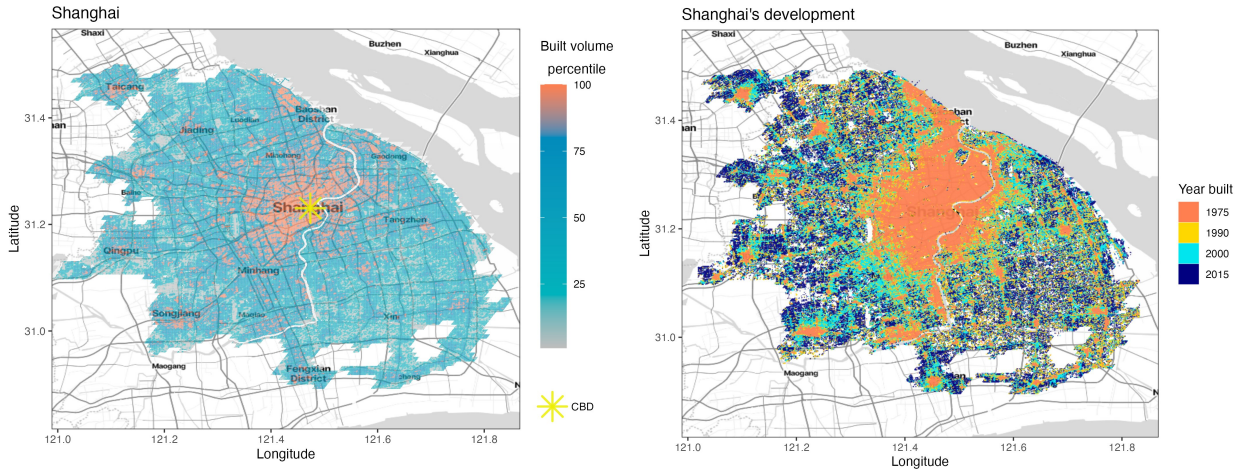


Figure 7: Visualization of the data used for estimation, for Shanghai. Left: percentiles of built volume at the 100m×100m resolution in Shanghai, 2015, with the CBD marked in yellow. Right: Shanghai's development over time.

Finally, in a second step, I shrink estimates $\widehat{\frac{\tau_i \gamma_i}{\beta}}$ towards the country mean using an empirical Bayes' procedure to enforce that $\widehat{\frac{\tau_i \gamma_i}{\beta}} > 0$; see Appendix D.1. The reason that this is necessary is that for model quantification, it must be that $\tau_i > 0$. However around 6% of the estimated skylines slope upwards. Upwards sloping skylines are estimated with considerably less precision than the vast majority of estimates (see the left panel of Appendix Figure A11) suggesting that noisy data drives these estimates, which the empirical Bayes procedure shrinks away.

4.3 Estimating the floorspace supply elasticity

To estimate the floorspace supply elasticity γ_i, I take logs of equation (14). Rearranging,

$$\log \frac{H_i}{\pi X_i^2} = \frac{\gamma_i}{1+\gamma_i}\left(\log w_i + \log \frac{L_i}{\pi X_i^2}\right) + a_i. \tag{18}$$

I use this as an estimating question for γ_i where city specific constants have been absorbed in a_i, which I treat as an error term. The left panel of Figure 8 shows this relationship in the data, and reveals that a log-linear relationship is a good approximation of the relationship between average height and income per square kilometer. However, this log-linear relationship is not directly informative of the floorspace supply elasticity.

First, the error term a_i contains $\log Z_i^H$.[33] Cities that have high productivity in floorspace construction will on average have lower floorspace prices, and in equilibrium, such cities will be larger in terms of L_i and πX_i^2, which can influence wages through price and agglomeration effects. Consequently, the regressors of interest are endogenous. Second, I only observe built volume in the cross-section, and consequently cannot estimate a separate γ_i for each city.

To handle endogeneity, I use instrumental variables. In particular, I instrument for $\log w_i$, which is measured in the data as city GDP/capita, using $\log \bar{Z}_i^y$, which I recover by inverting the calibrated model (see Section 5.1) and I control for population density, geophysical covariates G_i, and their interactions. Using this instrument isolates the component of wages driven by city productivity from variation in wages driven by goods prices and agglomeration effects. I include country fixed effects (ξ_n) to difference out country-specific determinants of height and productivity, so the orthogonality assumption is that $\log \bar{Z}_i^y \perp a_i \mid L_i/\pi X_i^2, \xi_n, G_i$. Put simply, the assumption is that tradeable goods productivity is uncorrelated with floorspace construction productivity within a country, conditional on population density and geophysical controls. For example, Los Angeles is productive because its fair weather and diverse landscape made it ideal for film production, but skyscraper construction is stymied by seismic activity and varied bedrock depth across the Los Angeles basin.

Second, to estimate city-specific floorspace supply elasticities, I parameterize $\widehat{\frac{\gamma_i}{1+\gamma_i}} = G_i\hat{\Gamma}$. G_i is a vector of city-level covariates, and Γ is a vector of coefficients. In particular, I include in G_i soil characteristics (whether there is clay, water, or sand in the soil and the soil's density), a city's elevation, average slope, and the fraction of warehouse development costs allocated to construction permits from the World Bank's 'Doing Business' survey, which is available at country-level averages, alongside city-specific measures for a handful of global cities. I form instruments for these interactions by interacting $\log \bar{Z}_i^y$ with the elements of G_i and estimate with two stage least squares (TSLS). Appendix D details the variables in G_i and the construction of these instruments.

[33] The error term is,

$$a_i = \frac{1}{1+\gamma_i}\log\left(\frac{Z_i^H}{\psi^H}\right) + \log\left(\frac{\left(1-\frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)^{\frac{\gamma_i}{1+\gamma_i}}}{1-\frac{\tau_i}{2}\frac{\gamma_i}{\beta}}\right).$$

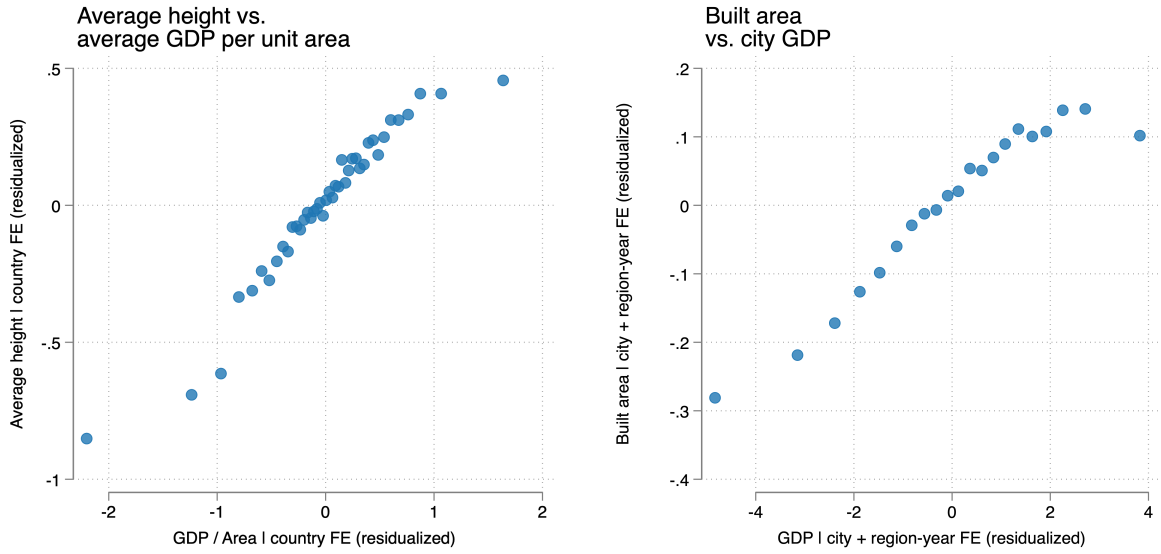


Figure 8: Conditional binscatters (Cattaneo et al., 2024) of the model-implied relationships between income and development and in the data. Left: average height versus GDP per square kilometer conditional on country fixed effects. Right: built area versus GDP, conditional on city and region-year fixed effects. The city fixed effects absorb a substantial amount of variation in the data, which is why the optimal number of bins is fewer in the right panel.

4.4 Estimating the land supply elasticity

Taking logs of equation (15) gives

$$\log \pi X_i^2 = \frac{\rho_i}{1+\rho_i} \log w_i L_i + b_i,$$

where b_i is a city-specific constant that contains $\log Z_i^X$.[34] I use this relationship as a basis to estimate ρ_i. It says that higher output cities are larger in terms of area. However, when productivity in urban land production, Z_i^X, is high, floorspace is less dear, and the city is more attractive. As a result, a naive regression of area on GDP would fail to identify ρ_i due to endogeneity bias.

To estimate ρ_i, I leverage panel variation in urban area and GDP, which is absent for built volume data. The panel variation here is compelling in that it uses how cities actually grow, rather than the cross-sectional

[34] The error term is,

$$b_i = \log \pi + \log \frac{Z_i^X}{\psi_n^X},$$

where I assume the quality-adjustment factor ψ_n^X varies across countries but not across cities within them. These are then netted out with a country fixed effect, alongside the constant, $\log \pi$.

relationship (which represents a long-run supply elasticity) and allows me to control for city and region-year fixed effects, which absorb unobserved shifters of both income and urban land construction that are common within cities over time or common across cities within a region-year. Using these fixed effects instead of productivity instruments relaxes the assumption that productivity in producing goods is orthogonal to productivity in producing floorspace or land. In particular, I estimate,

$$\log \text{Area}_{it} = \frac{\rho_i}{1+\rho_i} \log GDP_{it} + \xi_i + \xi_{rt} + e_{it}, \tag{19}$$

where ξ_i is a city fixed effect, ξ_{rt} are subcontinent-year fixed effects, and $t = 1990, 2000, 2015$. The right panel of Figure 7 displays this data for Shanghai, while the right panel of Figure 8 shows this time-series variation, and confirms the log-linear approximation fits the data well for all cities.

Again, I parameterize the land supply elasticity as a function of observables, $\widehat{\frac{\rho_i}{1+\rho_i}} = G_i\hat{\Omega}$. The identification assumption is that $\log GDP_{it} \perp e_{it} \mid \xi_i, \xi_{rt}$. The city fixed effect differences out all time-invariant components of urban land production productivity, which may be driven by fixed, geological factors. Urban land production productivity may be varying over time as technological advances in, e.g., electrical grid, paving, or sewage technology, have allowed cities to expand horizontally. When such trends are common at the sub-continent level, estimating (19) identifies ρ_i.

For GDP data, I am unable to construct a time series using the high-resolution nightlights (VIIRS) as the VIIRS satellite was launched in late 2011. Instead, I use the time series provided by the UCDB from Kummu et al. (2018), which is a population-weighted downscaling of subnational GDP measures. This GDP data is likely measured with substantial error, which would bias OLS estimates downwards. I assume this measurement error is classical and use the DMSP-OLS nightlights time series to instrument for city GDP and its interactions with G_i to handle this form of attenuation bias.[35]

Sample restrictions For estimation, I restrict the sample to cities in countries where urban form more likely reflects market forces. To that end, I omit former and current communist nations, including East Germany, as in these nations, urban planning was or continues to be state-directed. Second, I drop nations' primate

[35] An alternative would be to directly impute city income with the DMSP-OLS data. However, the DMSP-OLS data has substantially more top and bottom coding than the VIIRS data, and so there is no guarantee that income measures constructed with it would have less measurement error than that from Kummu et al. (2018).

(largest) cities, given that in the developing world, political factors often explain their size (Ades and Glaeser, 1995), and consequently the correlation between income, height, and area might be marred by an omitted variable (weak provision of public services elsewhere or dictatorial command-and-control).

4.5 Estimation results

Floorspace supply elasticity estimates Regression estimates are available in Appendix Table A10. I estimate γ_i to be on average 0.95, with substantial heterogeneity across countries and cities within countries. The left panel of Figure 9 shows the relationship between a country's average γ_i and its income per capita. Across all cities, 90% of estimates are between 0.33 and 1.68. Slope, elevation, and sand and clay content predict lower-than average floorspace supply elasticities, while denser soil and increased water content in soil predict larger elasticities. Across countries, a larger regulatory burden in construction is associated with lower-than-average floorspace supply elasticities. A city's floorspace supply elasticity is on average increasing in a nation's GDP/capita, though I estimate that floorspace supply is substantially more elastic in Japan (on average, 1.50) compared to the United States (1.01). This accords with narrative evidence that among rich-world countries: Japan is not facing a housing crisis driven by self-imposed housing supply restrictions.[36] Cities in very poor nations like those in sub-Saharan Africa have floorspace supply elasticities around 0.45, while floorspace is more elastically supplied in nations like India and China. The top panel of Figure A2 displays a map of γ_i across all cities. East Asia, some parts of India, and southern Europe stand out has having particularly elastic floorspace supply. Within countries, there is dispersion in estimates. For example in China, floorspace appears less elastically supplied in the smaller inland cities compared to its urban giants.

Commuting costs elasticity estimates To recover τ_i, I divide the estimated skyline slopes $\widehat{\tau_i\gamma_i/\beta}$ by $\hat{\gamma}_i/\beta$ where I assume $\beta = 0.25$. This value accords with estimates from the literature; for example, Davis and Ortalo-Magné (2011) estimate this at 0.24 for the United States. It is likely that this value varies across countries. Using consumption expenditure data from the World Bank's 2017 International Comparison Program (ICP), I show in the left panel of Appendix Figure A1 that the expenditure share of housing is rising in a nation's GDP/capita: poorer countries have lower β, which would flatten the observed relationship between

[36] See, e.g., Appelbaum, B. "The Big City Where Housing Is Still Affordable." (2023, September 11). *The New York Times.*, or "The growing global movement to restrain house prices." (2023, September 6). *The Economist.*

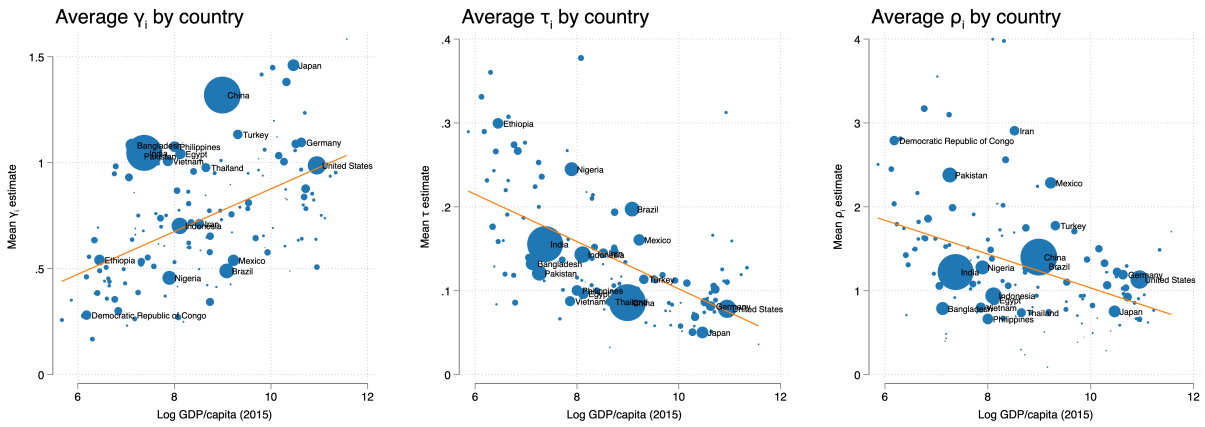


Figure 9: Simple averages across cities within a country of floorspace supply elasticities (γ_i, left), commuting costs (τ_i, middle) and land supply elasticities (ρ_i, right) versus a country's log GDP per capita.

income level and urban commuting costs. Using β estimates that vary by country attenuates the relationship visible in Figure 9 but the pattern that commuting costs are higher in developing nations persists; see the right panel of Figure A1. In other literature, most income elasticity estimates for housing are one or less (see the discussion in Section 3), suggesting that β should fall in a nation's GDP per capita. This discrepancy may arise as poorer nations have larger rural populations that may record little housing expenditure on such surveys. As my focus is on technological differences in cities around the world, I hold this value fixed across countries.

Estimates of τ_i holding β fixed are on average 0.14. East Asian nations have some of the lowest estimated commuting costs with values between 0.05-0.06 for Japan, South Korea, and Taiwan.[37] OECD nations have values of τ ranging from 0.08 (USA) to 0.12 (Chile). τ_i is on average quite high for cities in the developing world, with cities in sub-Saharan Africa having a τ_i on average 0.25, with heterogeneity across the sub-continent: in Uganda the average is 0.27, while in comparably rich Botswana, it is 0.11. The middle panel of Figure 9 shows that τ_i is on average falling in a nation's GDP per capita.

There is noticeable dispersion of τ_i across cities within a country, as shown in the middle panel of Appendix Figure A2. Looking within countries, larger cities have on average lower τ_i, suggesting that τ_i indeed captures a negative amenity: cities with worse transportation infrastructure are smaller. Appendix Figure A4 plots

[37] In former Soviet nations, which are excluded from the analysis sample, I recover comparable estimates of τ_i. These range from from 0.07 (Latvia) to 0.12 (Russia). Among these countries include authoritarian states that to this day significantly control urban development, like Turkmenistan (0.09), Uzbekistan (0.09), Kazakhstan (0.11), and Tajikistan (0.12), which suggests that the estimation is indeed recovering differences in urban form, whether market-drive or state-planned.

the relationship between log city population and τ_i conditional on country fixed effects, and shows a clear downwards-sloping relationship. Within countries, a 1% increase in τ_i is associated with a 0.6% decrease in population (p-value < 0.01) on average. This relationship could be, of course, endogenous, as larger cities invest more in transportation infrastructure because they have the funds or incentive to do so. In this paper, I view τ_i as a city's technology, though accounting for endogenous congestion or the political economy of urban infrastructure investment is an active research area (Allen and Arkolakis, 2022; Bordeu, 2023).

Land supply elasticities Across cities globally, I estimate that the average land supply elasticity, ρ_i, is 1.46 and ranges from 0.09 to 17.17. Richer countries tend to have less elastic land supply on average (right panel of Figure 9). This average may seem surprisingly low given how much undeveloped land there is in most parts of the world. The majority of habitable land is neither uninhabited nor used for agriculture, and urban land represents a tiny fraction of overall habitable land (Ritchie and Roser, 2019). This might suggest that land is in effectively free supply. However, even in the 2015 cross section, the OLS elasticity of area to GDP is 0.65 (conditional on country fixed effects, s.e. 0.007), implying an average land supply elasticity of 1.85. This cross-sectional relationship represents a long-run land supply estimate (rather than one estimated in my 'medium-run' 10-15 year intervals), and is marred by endogeneity concerns that bias the estimate upwards (cities with higher productivity in urban land production are both larger in terms of income and area). Thus, even in the long-run, and absent endogeneity concerns, urban land is far from perfectly elastically supplied.

Using my time series estimates, I find that the largest driver of differences in land supply is the share of developable land within 10km of the CBD. Elevation and slope are positively correlated with land supply, while poor soil content is a negative predictor of elastic land supply, suggesting that building out is a substitute for building up. The land supply elasticity is fairly low in East and Southeast Asian countries (on average, 0.87), many of which are islands or predominately coastal. In land-abundant sub-Saharan Africa, the average land supply elasticity is 2.09. In Europe and North America, the average land supply elasticity is 1.13, which is the average value in the United States, compared to 0.66 for land-constrained Denmark. These estimates are dispersed across cities within countries; see the bottom panel of Appendix Figure A2. For example, some cites are very land-constrained in Brazil, while others face fairly elastic land supply. This dispersion stands in contrast to an island nation like Indonesia, where most cities are equivalently land-constrained.

Urban cost elasticities Summarizing the estimates so far, cities in poor nations grow by building *out*, not *up*.

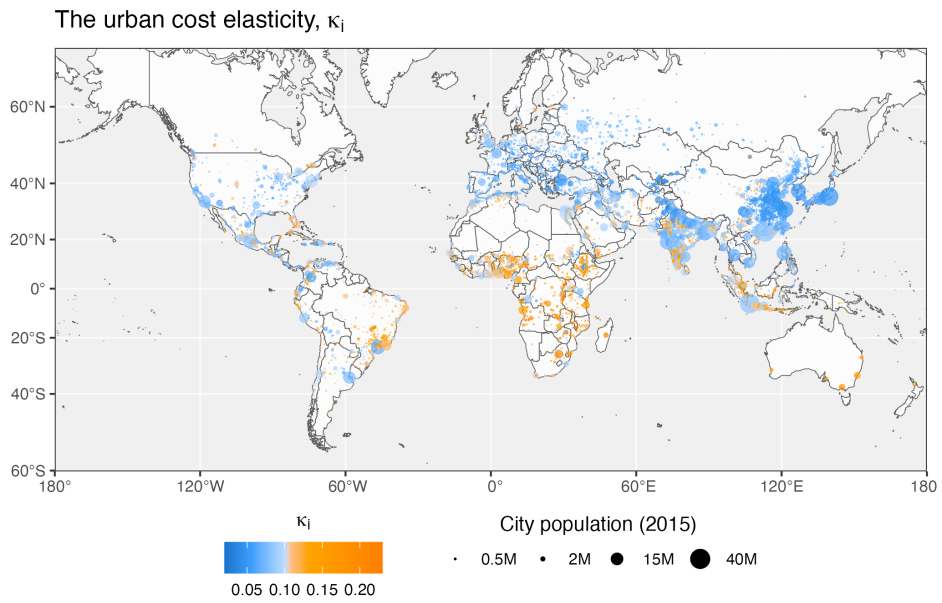


Figure 10: The spatial distribution of cities, colored by κ_i and sized according to their population.

While this greater ability to expand horizontally ameliorates congestion forces brought about by competition for floorspace on extant urban land, it means households in poor nation cities must on average travel farther to commute to work. As cities' commuting costs are higher in the developing world than those in the developed world, this suggests that urban costs may be higher in poor nation cities.

To measure the urban cost elasticity, I construct κ_i for each city using equation (11) and the estimated parameters of its urban technology, τ_i, γ_i, and ρ_i. Across cities globally, the average urban cost elasticity is 0.09. In rich nations like France or the United States, the average urban cost elasticity is around 0.08. Combes et al. (2019), using a similar framework and data on house prices, estimates κ_i to be on average 0.03 in France. The urban cost elasticity is on average higher in poor nations. In nations with an GDP/capita $<$ 4000 USD, the urban cost elasticity is on average 0.11, which is 35% higher than it is in high-income nations. In sub-Saharan Africa, it is on average 0.15. These differences are visually apparent. In Figure 10, I map cities' urban cost elasticities around the entire world. Cities shaded orange have higher urban cost elasticities than average and are predominately located in developing nations, while developed nations, China, and former Soviet nations have cities with lower urban cost elasticities on average. Appendix Table A2 summarizes the average urban cost elasticity and its components by region. The left panel of Appendix Figure A3 plots the average urban cost elasticity across cities within a country against a nation's log GDP per capita. The

relationship is strongly negative.

The biggest driver of differences in the urban cost elasticity across cities is τ_i. In Appendix Table A3, I report regression coefficients and partial R^2 statistics from a regression of κ_i on τ_i, γ_i, and ρ_i. The coefficient on τ_i is an order of magnitude larger than coefficients on γ_i and ρ_i, and its partial R^2 is 0.47, while it is 0.24 and 0.02 for γ_i and ρ_i, respectively. 45% of the variation in κ_i is across countries rather than across cities. In Appendix Table A5 I report results of a regression of κ_i (in levels and logs) against a nation's GDP/capita as well as against a city's GDP/capita and population, controlling for country fixed effects. Coefficients on a country's log GDP per capita are over twice as large as they are for city level GDP per capita: doubling a nation's income per capita is associated with a 9% decline in the urban cost elasticity. Looking within countries, richer and larger cities have lower urban cost elasticities – doubling income is associated with a 4% decline in the urban cost elasticity. This accords with the logic the theory: these cities can absorb population without putting upward pressure on urban costs, which allows them to be larger in equilibrium.

Levels of urban costs Using my estimates of γ_i, τ_i, and ρ_i, alongside data on city area, total built volume, and city GDP, I can recover values of Z_i^H/ψ_n^H and Z_i^X/ψ_n^X for each city i in nation n. To separate the quality adjustment terms ψ_n^H and ψ_n^X from physical productivity, I use data on average floorspace prices across nations. In aggregate, the model implies,

$$\psi_n^H q_n \sum_{i=1}^{N_n} H_i = \beta \sum_{i=1}^{N_n} w_i L_i \quad \text{and} \quad \psi^X q_n \sum_{i=1}^{N_n} \pi X_i^2 = \sum_{i=1}^{N_n} \frac{\beta}{1+\gamma_i} w_i L_i$$

where q_n is the average price of urban floorspace in country n. For price data, I use data from Liotta et al. (2022). Their data measures floorspace prices per square meter in local currency for some cities in 49 countries (often, only the largest city), including developing world countries like Ethiopia, Cote d'Ivoire, Indonesia, and Pakistan. I convert their estimates to USD equivalents using the World Bank's PPP deflator. For countries outside their sample, I predict the average floorspace price using a Random Forest and use the PPP deflator, the World Bank's measure of the local regulatory burden, a nation's log GDP/capita, and its size measured in population as predictors. I use the actual and predicted q_n values to recover ψ^H and ψ^X and net them out of the recovered values Z_i^H/ψ_n^H and Z_i^X/ψ_n^X.

To compare levels of urban technology across countries, I compare differences in the product $Z_i^H Z_i^X$ across

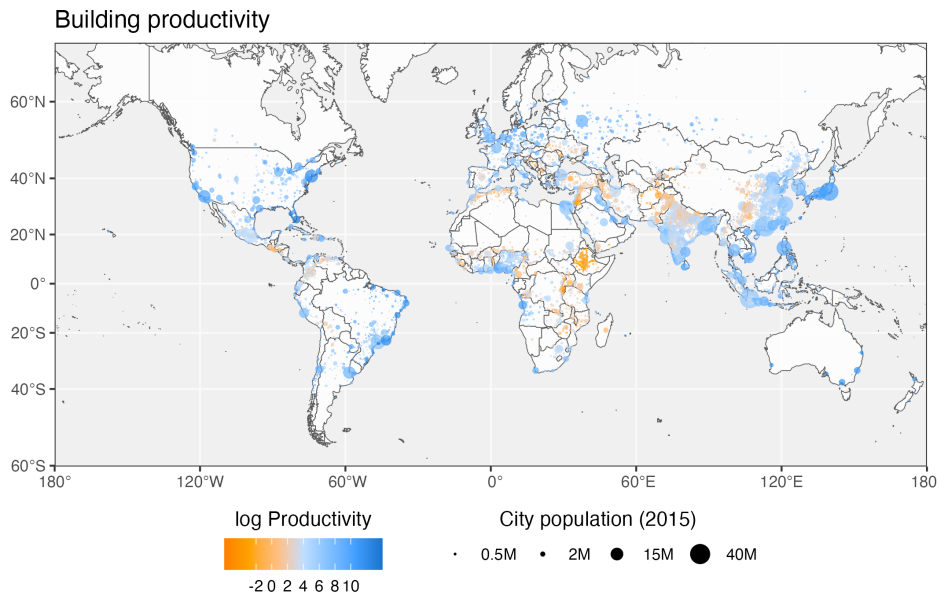


Figure 11: A map displaying variation in building productivity across cities globally. The correlation with average levels of building productivity versus a nation's log GDP/capita is available in the right panel of Appendix Figure A3.

cities, which I refer to as building productivity. I use this aggregation because vertical productivity (Z_i^H) and horizontal productivity (Z_i^X) in urban development enter the utility function multiplicatively, albeit raised to different powers. Figure 11 displays the spatial distribution of building productivity across cities globally. Building productivity is low in sub-Saharan Africa, India, and Southern and Eastern Europe. However, within countries, the largest cities often have a higher level of building productivity than smaller cities. For example, Lagos, Nigeria, has a high level of building productivity compared to the national average. Globally, among major cities, some of the highest building productivity cities include Adelaide, Australia, Jeddah, Saudi Arabia, New York, United States, and Rio de Janeiro, Brazil. On average, building productivity is 3.6 log points higher in developed nations' cities, compared to those in the developing world, a massive gap. The correlation between the average log building productivity across cities and a nation's log GDP/capita is given in the right panel of Appendix Figure A3 and is strongly positive. Building productivity is negatively correlated with the urban cost elasticity; Appendix Figure A5 displays the relationship between a city's urban cost elasticity and its building productivity, conditional on country fixed effects. The relationship is on average downward sloping though it displays a somewhat U-shaped relationship, where the cities with the highest building productivity on average have slightly higher urban cost elasticities.

Model validation and comparison of estimates Appendix E details many model validation checks in a context where data is abundant: the United States. First, I confirm the model's prediction on the structure of cities aligns with predictions of the monocentric model. The residency gradient declines in distance to a city's CBD, floorspace prices fall with respect to distance to the CBD, and most employment is concentrated in the CBD. I then ask whether my parameters can predict city structure. τ_i strongly predicts the price gradient and the population gradient, while γ_i only predicts U.S. cities' population gradient, consistent with the theory. U.S. cities with higher values of τ_i, γ_i, and ρ_i have lower average floorspace prices, consistent with τ_i being a negative amenity reflected in prices, and γ_i and ρ_i controlling the total supply of floorspace in a city. My measures of γ_i are weakly positively correlated to those estimated in Baum-Snow and Han (2024), and my estimates of ρ_i are strongly positively correlated with similar estimates from Saiz (2010).

There is no other paper that computes τ_i at the city level as I have done here. However, Akbar et al. (2023b) compute measures of freeflow and daytime speed in cities and their downtowns for over 1,000 global cities. While their definition of urban agglomeration is not identical to mine, I am able to match 88% of their cities to those in my sample, though I drop 117 of these cities that are in countries with a history of state-driven urban planning.[38] Appendix Table A4 shows correlations of log τ_i with the log of measured speed near city center and the speed indices computed in that paper. Holding midday speed fixed, cities with faster speeds downtown at midnight have substantially lower τ_i, and this relationship is statistically significant ($p < 0.05$), and robust when controlling for city characteristics and country fixed effects. Interestingly, faster cities at midday have higher τ_i, suggesting that low τ_i cities may be endogenously more congested (columns 1-4). These results broadly accord with using the paper's uncongested and daytime speed indices in place of measures of speed in the city center (columns 5-8).

Using data from the OpenStreetMap project (Boeing, 2024), I correlate τ_i with characteristics of each city's road network. Consistent with the hypothesis that cities with better transportation infrastructure have lower τ_i, estimates of τ_i are negatively correlated with the share of roads paved and the total amount of road in a city. In Appendix Table A7, I regress the log of τ_i on the share of roads paved, log total length of streets in a city, the log average number of lanes per kilometer, the log number of intersections, and the log average

[38] The GHSL urban agglomerations are on average larger than those in the Akbar et al. (2023b) data. For example their data does distinguish Singapore and Johor Bahru, which are split by the Malaysian border, but appear a single urban agglomeration in GHSL. I associate urban agglomerations that cross-national borders with the nation where most of the city lies. Within countries, many cities defined by political boundaries form a single urban agglomeration in the GHSL data, e.g., San Francisco and San Jose in the United States.

road width, and control for the log population of a city and country fixed effects, as well as an indicator for whether a city is its nation's largest. I find that a 10pp increase in the share of roads paved is associated with a 3.3% reduction in τ_i ($p < 0.01$), and that doubling the amount of available roadway in a city lowers τ_i by 14% ($p < 0.1$), though the additional transportation network characteristics have no statistically significant correlation with τ_i.

5 How important are urban costs?

In this section, I ask how urban costs affect a nation's level of development, what policies are available to cost-effectively lower these barriers to urban development, and how urban costs govern a nation's capacity to adapt to urbanizing shocks. Urban costs reflect both geophysical constraints and policy variables that are unmodeled but nonetheless are reflected in my parameter estimates. Thus, my parameter estimates are not robust to the Lucas critique, and can potentially change endogenously as cities invest in transportation infrastructure or reduce regulations that inhibit both vertical and horizontal urban growth. I nonetheless can answer several related questions: what would a nation's urban system look like if its technology were brought closer to the frontier? What policies might be associated with reducing urban costs, and are such policies cost effective? How do economies with different urban systems react to the same shock?

To answer these questions, I perform two counterfactuals. First, I ask what the welfare gains are associated with lowering urban costs in developing nations. To do so, I lower the level of urban cost elasticities within a country so that the country average level is equal to the average level observed in the United States. This is not a policy counterfactual. Instead, my aim is to get a sense of the stakes, and ask how differences in urban costs contribute to nations' overall level of development. Second, to study feasible policy, I ask what variables might affect urban costs in the developing world. To that end, I study road paving. Urban roads in developing nations are substantially under-paved compared to those in rich nations. The share of roads paved is correlated with my measure of urban costs, and is a policy whose costs I can quantify.

Parameter	Value	Description	Source
ζ	0.04	Elasticity of urban productivity with respect to density	Combes et al. (2010) and Ahlfeldt and Pietrostefani (2019)
σ	4	intercity trade elasticity	Bajzik et al. (2020)
β	0.25	share of income spent on housing	Average across countries where observed (World Bank 2017 ICP)
$1-\alpha_n-\beta$	–	Share of income spent on agricultural goods	Calibrated to match World Bank Development Indicators in 2015 on the share of agricultural value-added in national income
μ	0.7	share of land in agricultural goods production	Chari et al. (2021)
ε	1.17	migration elasticity	Suárez Serrato and Zidar (2016) and Sahai and Bailey (2022)

Table 1: Externally calibrated parameters. The parameter α_n varies across nations n so aggregate expenditure on agricultural goods in the model matches the share of value-added attributed to the agricultural sector in the data.

5.1 Model quantitification

To perform this counterfactual analysis, I need to quantify the rest of the model. Quantification follows three steps. First, in addition to the elasticities estimated in Section 4, I select values for the remaining model elasticities from the literature. I then compute trade costs δ_{ij} across all cities globally. Finally, I invert the model by recovering fundamentals (e.g., A_i) that exactly rationalize the data in 2015 as a model equilibrium.

Externally calibrated parameters Table 1 reports all externally calibrated parameter values. For the elasticity of urban TFP with respect to density, I set $\zeta = 0.04$, which is the average wage elasticity with respect to density in the meta-analysis of Ahlfeldt and Pietrostefani (2019) and is the elasticity of urban TFP with respect to density in Combes et al. (2010). There is some evidence that agglomeration economies may be stronger in the developing world (Chauvin et al., 2017), but the evidence is not conclusive. Estimating agglomeration economies is notoriously difficult (Combes et al., 2011) and not the focus of this study, so I choose the same parameter for all cities. I set the elasticity of substitution across urban varieties $\sigma = 4$, which is near the median value across estimates of this elasticity (across countries) in Bajzik et al. (2020) and is used in similar studies.[39] For the land share in agriculture, I set $\mu = 0.7$ which is around the average of land shares in agricultural production estimated in Chari et al. (2021) in China for major crops like

[39]For example, Monte et al. (2018) use $\sigma = 4$ for the elasticity of substitution across U.S. county varieties.

wheat, rice, and sugar.[40] For the migration elasticity, I choose $\varepsilon = 1.17$. This value is on the low-end for the literature and is intentionally chosen so that my estimates of the gains of reallocation of population in counterfactual analysis are conservative. This value comes from Sahai and Bailey (2022), who use migration data from Facebook and temperature shocks to identify the migration elasticity in India. Using business tax shocks in the United States, Suárez Serrato and Zidar (2016) estimate an interstate migration elasticity of 1.20.[41] I choose $\beta = 0.25$, which is close to the average housing expenditure share in the World Bank 2017 ICP survey in sampled countries. I choose α_n to match the residual income share that is neither agriculture or housing for each nation n, using the 2015 World Development Indicators data for agriculture's share of value-added.

Trade costs In the model, each country is in autarky, but within countries, cities can trade with each other and the agricultural sector, subject to trade costs. The agricultural good is the numeraire and freely traded. To get trade costs across cities, I first estimate the average drive time between cities in every country using the Open Street Maps routing data, d_{ij}. To convert this to iceberg trade costs, I assume $\delta_{ij} = (1 + d_{ij})^{-\bar{\delta}}$ where d_{ij} is measured in kilometers.[42] Trade costs only enter the model raised to the power $1 - \sigma$, so I seek to estimate $-\bar{\delta}(\sigma - 1)$. To that end, I estimate a gravity regression with a pseudo-Poisson maximum likelihood estimator (Silva and Tenreyro, 2006) by projecting the value of road-based shipments across cities in the United States onto log distance, controlling for origin and destination fixed effects, using the 2017 Commodity Flow Survey (CFS). Appendix Table A6 reports the results of this regression; I obtain a coefficient $\widehat{\bar{\delta}(\sigma - 1)} = 0.92$, which is close to the value obtained using the CFS microdata in Dingel (2017).

Armed with these parameters, I invert the model for each country to find values of $\bar{Z}_i^y$, A_i, Z_i^H, and Z_i^X which exactly rationalize data on city-level GDP, population, agricultural output, and city size (measured in km^2) as an equilibrium of the model.

[40] Their average estimate across crops is 0.47, but this includes varieties like fruits and vegetables that are not a large share of agricultural output globally and are not very land-intensive.

[41] On the higher end, in Indonesia, Bryan and Morten (2019) find a substantially higher migration elasticity of 3.2, but this estimate stems from a model in which migration occurs through the spatial sorting of workers with heterogeneous skills, and reflects skill dispersion.

[42] For trade costs of urban varieties to the agricultural sector, I assign the average exporting cost faced by a city, so that $\delta_{ia} = \frac{1}{I_n - 1} \sum_{j \neq i} \delta_{ij}$.

5.2 Lowering urban costs to the U.S. level

The first counterfactual I consider is lowering the urban costs elasticity to the level observed in the United States, which is close to the technological frontier. The object of interest is a nation n's welfare, $\mathcal{W}_n$. Welfare is the ex-ante expected utility of residing in any city or the rural sector, and is given by,

$$\mathcal{W}_n = \left(\sum_{i=a,1,..,I_n} \left(A_i (Z_i^H)^{\frac{\beta}{1+\gamma_i}} (Z_i^X)^{\frac{\beta}{1+\gamma_i} - \frac{\tau_i}{2}} \frac{w_i}{P_i^{\alpha_n}} (\beta w_i L_i)^{-\kappa_i} \right)^{\varepsilon} \right)^{1/\varepsilon}. \tag{20}$$

Defining $\hat{A}_i = A_i (Z_i^H)^{\frac{\beta}{1+\gamma_i}} (Z_i^X)^{\frac{\beta}{1+\gamma_i} - \frac{\tau_i}{2}}$, to the first order, any change in the level of urban costs (Z_i^H, Z_i^X) or the urban cost elasticity (κ_i) is given by,

$$\frac{d\mathcal{W}_n}{\mathcal{W}_n} = \sum_i \left(\frac{L_i}{L_n} \right) \left(\underbrace{\frac{d\hat{A}_i}{\hat{A}_i}}_{\text{direct level effect}} + \underbrace{\frac{d(w_i/P_i^{\alpha_n})}{(w_i/P_i^{\alpha_n})} - \kappa_i \frac{d(w_i L_i)}{w_i L_i}}_{\text{real wage effect}} + \underbrace{\kappa_i \log(\beta w_i L_i) \frac{d\kappa_i}{\kappa_i}}_{\text{direct urban cost effect}} \right). \tag{21}$$

Equation (21) decomposes to the first order the mechanisms by which lowering urban costs affects welfare. First, by lowering urban costs, there is a direct effect on the level of utility in each location, by shifting the floorspace supply curve, captured by $d\hat{A}_i/\hat{A}_i$. The second effect is what I refer to as the real wage effect, and captures how value of income spent on traded goods adjusts in equilibrium when urban costs change. Absent agglomeration forces, the model would be efficient, so by the envelope theorem, changes in fundamentals would capture to the first order changes in welfare, and these real wage effects (i.e., price effects) would be zero. However, in this counterfactual, I allow there to be agglomeration forces, and consider large changes to urban costs, which leads to a substantial labor reallocation and large changes in real wages across space. The average real wage change across cities capitalizes the gain from the spatial reallocation of labor and therefore measures the component of the welfare gain attributable to general equilibrium mechanisms. When workers sort to more productive cities, wages may rise due to the endogenous productivity response through agglomeration. Moreover, by sorting to more productive cities, the average price of traded goods falls, which is measured by P_i. The real wage effect in the agricultural sector captures the general equilibrium gains from rural-to-urban migration. As workers leave the agricultural sector to cities, the agricultural wage rises. Rural-to-urban migration continues until the marginal migrant is indifferent between the rural and urban sectors,

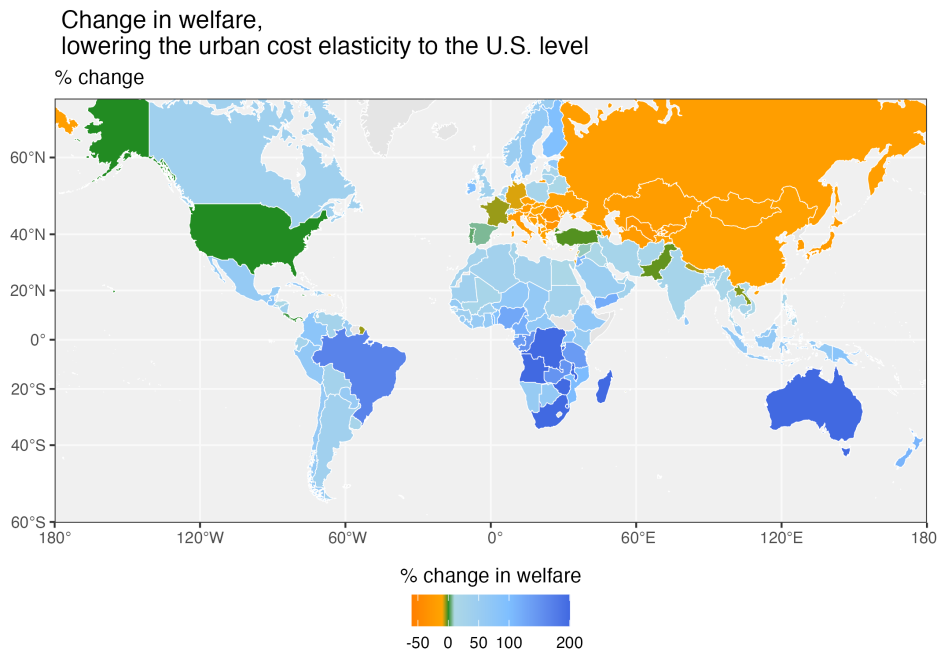


Figure 12: Welfare changes (% terms) from changing the urban cost elasticity (κ_i). Values are capped at 200%. While welfare rises for the average person by only 40%, and 95% of individuals experience welfare changes between -56 and 169%, welfare gains are higher in several regions of sub-Saharan Africa: for example, in Angola, welfare increases by 338%, or in the Democratic Republic of the Congo, where welfare rises by 568%.

so the rise in agricultural real wages capitalizes the gain associated with urbanization. Finally, the third term captures the urban cost effect. The third term measures the direct effect associated with rotating the 'urban cost curve' by changing κ_i. This decomposition allows me to disentangle the direct effects of lowering urban costs from the general equilibrium effects due to labor reallocation.

To isolate the role of changing the urban cost elasticity (κ_i), I shock κ_i, holding $\hat{A}_i$ at their baseline levels, and use the scale of labor and wages in the data. I solve for counterfactuals in levels by recomputing the full equilibrium under different parameter values for every country in the data. I change κ_i so that on average, it is equal to the U.S. value across cities in every country. Figure 12 maps the change in welfare across countries. From the map, it is visually apparent that the largest gains are in the developing world: welfare gains are concentrated in Africa, Latin America, and Southeast Asia. Nations like China, East Asian countries like Japan and South Korea, and former Soviet nations experience welfare losses, reflecting the fact that their urban technology is better than the U.S. level at baseline.

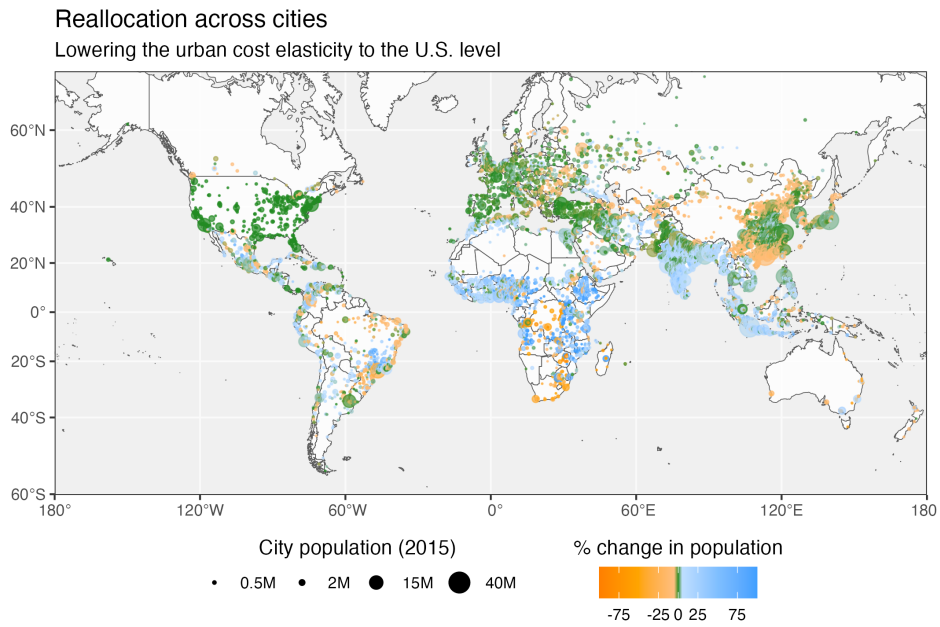


Figure 13: City population changes (% terms) from changing the urban cost elasticity (κ_i).

Welfare effects are large, because welfare measures the direct effect of lowering urban costs, which is mechanically positive for nations with higher urban cost elasticities than the U.S. on average. In developing nations (GDP/cap <$4,000), the population-weighted average welfare gain is 66%, compared to developed nations (GDP/cap > $20,000), where the gain is on average only 11%. Agglomeration effects play a small role in this story: holding productivity fixed at its baseline level, welfare gains in the developing world fall by only 1.4 percentage points. Appendix Figure A6 plots welfare and urbanization changes across nations against their baseline log GDP per capita. For developing nations, a sizable share of these welfare gains come from rural-to-urban migration. In developing nations, the share of the population living in cities rises by 8.5%, while in developed nations, the average change is less than 0.2%.

To disentangle the mechanical changes in welfare due to changing the urban cost elasticity from general equilibrium effects on aggregate productivity, I examine how real wages in the urban and agricultural sector respond to the changes in urban cost elasticities. Changes to the amount of wage income that can be spent on traded goods reflects changes to a nation's aggregate productivity and their overall level of development. These general equilibrium effects arise from workers moving to cities or from reallocating across them. Figure 13 shows how the spatial distribution of urban population changes around the world. In the developing world, many cities' populations grow, however there is noticeable within country heterogeneity, with some

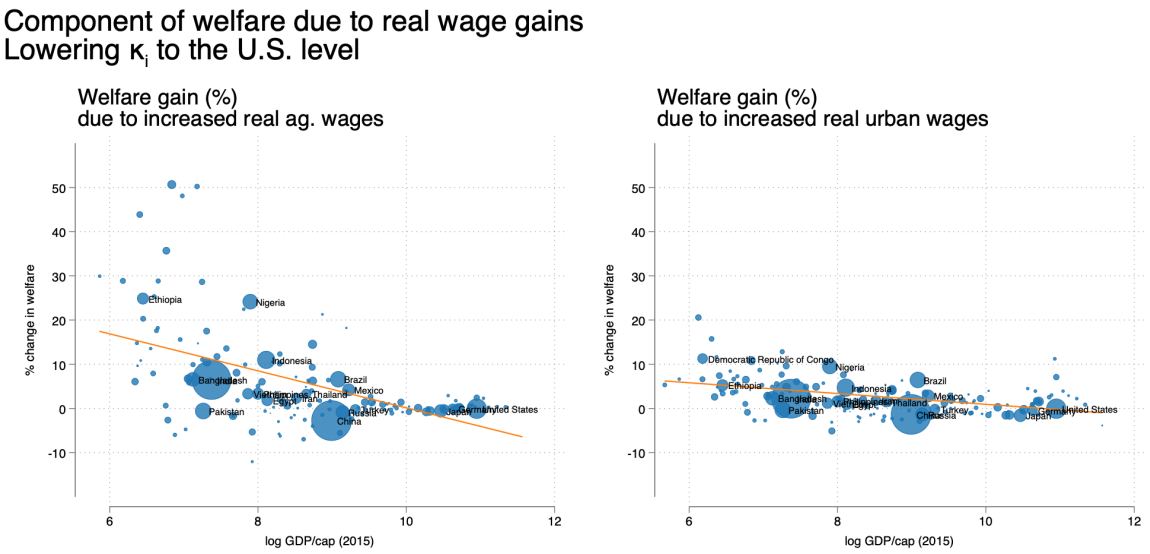


Figure 14: Welfare changes due to changes in aggregate productivity versus a nation's baseline log GDP/capita. Left: real wage changes in the agricultural sector. Right: average real wage changes in the urban sector.

cities shrinking and other growing, reflecting that the urban population reallocates across cities. When urban costs are lower, workers reallocate to higher-value cities that were previously not able to accommodate them.

Figure 14 plots the change in welfare due to changes in real agricultural and urban wages against a nation's baseline GDP per capita. In developing nations, the average, population-weighted real wage change in agriculture is 16%, while in cities it is 3.5%, implying that lowering the urban cost elasticity gives rise to an aggregate productivity increase of about 20%. However, in developed nations, aggregate productivity effects are small, reflecting the fact that these nations have a low share of their population working in agriculture, and suggesting that the baseline population distribution is not misallocated across cities.

Which drivers of the urban cost elasticity are most important? In Appendix Figure A7, I report welfare changes across regions from changing each component of the urban cost elasticity. Regardless of which component changes, the broad pattern stays the same: welfare effects are largest in the developing world, and are small or negative in developed nations, China, and former Soviet states. Consistent with the results reported in Section 4.5, the biggest driver of differences in urban costs across space is τ_i, the commuting cost elasticity. Changing just τ_i generates equivalent or larger gains than changing the average urban cost elasticity in most places, while changing just the floorspace supply elasticity, γ_i, increases welfare by

about half the magnitude of changing τ_i. Changing the land supply elasticity has negligible effects except in land-constrained Southeastern Asia, where reducing the land supply elasticity has larger welfare effects than changing τ_i. This is because cities in Southeastern Asia have lower commuting cost elasticities than those in other developing nations, so allowing cities to expand reduces the component of the urban cost elasticity driven by rising floorspace prices.

These exercises demonstrate the potential for lowering urban costs as a way to promote economic development and reduce international inequality. However, it does not reflect implementable policy, nor does it reflect the potential costs associated with improving cities. In the next section, I study one policy aimed at lowering urban costs in developing countries: paving cities' roads.

5.3 Paving roads to reduce urban costs

The quality of transportation infrastructure substantially varies across space. Again using data from the OpenStreetMap project, I show in Figure 15 that the share of urban roads (within 10km of the CBD) that are paved is strongly increasing in a nation's GDP/capita (left panel). The richest nations have almost all of their urban roads paved, but in poorest nations only somewhere between 25-50% of city roads are paved. Middle income countries have around 80% of their roads paved. Looking within countries, larger cities have a larger share of their roads paved (right panel). In short, the cities with the worst transportation infrastructure are small cities in the developing world.

Some attention in the development literature has focused on the impacts of road paving.[43] McIntosh et al. (2018) studies the impact of road paving in urban Mexico, finding that the return on investment on road paving, as measured by its capitalization in real estate values was around 100%, while Gonzalez-Navarro and Quintana-Domeque (2016) find a more modest 9% return on investment when studying road paving in a single city in Mexico.[44] I focus on the efficacy of road paving on lowering urban costs around the world. This

[43] The majority of this literature focuses on road paving to foster market access between cities and agricultural regions. For example, Storeygard (2016) studies how road surface affects transportation costs and the income of cities in sub-Saharan Africa, while Aggarwal (2018) studies the effect of paved roads on market integration in India villages. My framework allows me to assess the role of these intercity transportation costs on national income, but as the focus of this paper is on urban costs, I omit this analysis.

[44] Additionally, both Olken (2007) and Wong et al. (2013) study the efficacy of local governments in providing paved roads in Indonesia and China, respectively, but these studies focus on the role of corruption and local governance in effectively implementing this policy.

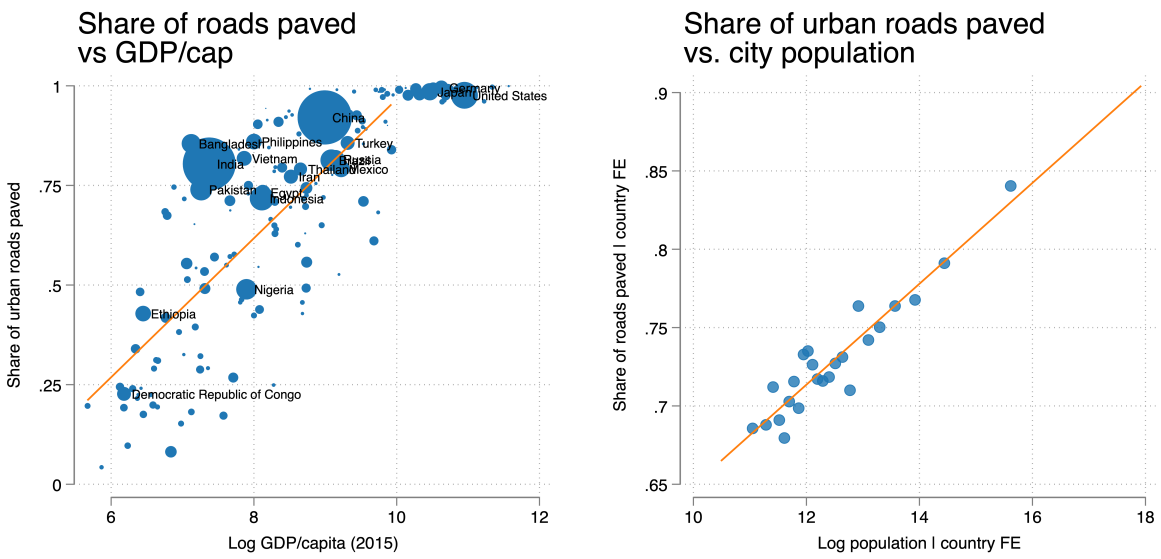


Figure 15: Road paving across cities and countries. Left: Avg. share of urban roads paved vs. national GDP/capita. Right: Share of roads paved vs. city population, conditional on country fixed effects

is a useful policy intervention to study, because the costs of road paving are well-understood. For example, the African Development Bank (2014) estimates the costs of construction and upgrading of paved roads to range from $115,900 to $425,400 per kilometer, depending on the scale of the project.

To assess the efficacy of road-paving, I first require a mapping from the share of roads paved to the quality of cities. To do so, I recover the relationship between road paving and τ_i, the commuting cost distance elasticity. Across cities, on average a 10pp increase in the share of roads paved is associated with an 8% decrease in the commuting cost elasticity (p-value <0.01). However, this baseline correlation likely does not reflect the causal effect of road paving on $\log \tau_i$. To isolate the causal effect, in Appendix Table A7, I regress the $\log \tau_i$ on the share of roads paved, controlling for additional transportation infrastructure variables, like the total log length of roadways in a city, the log population of a city, and an indicator for whether a city is a nation's largest, and country fixed effects. These covariates handle several issues with using the baseline correlation. Including additional transportation variables handles omitted variables bias: the share of roads paved may be correlated with other transportation variables that influence τ_i. Second, cities may be endogenously bigger because they have low τ_i, and thus they have more available funds to invest in road paving, so I control for a city's population. Additionally, as most of the variation in road-paving and τ_i is across countries and may reflect a variable related to a nation's overall level of development, I include country fixed effects. These

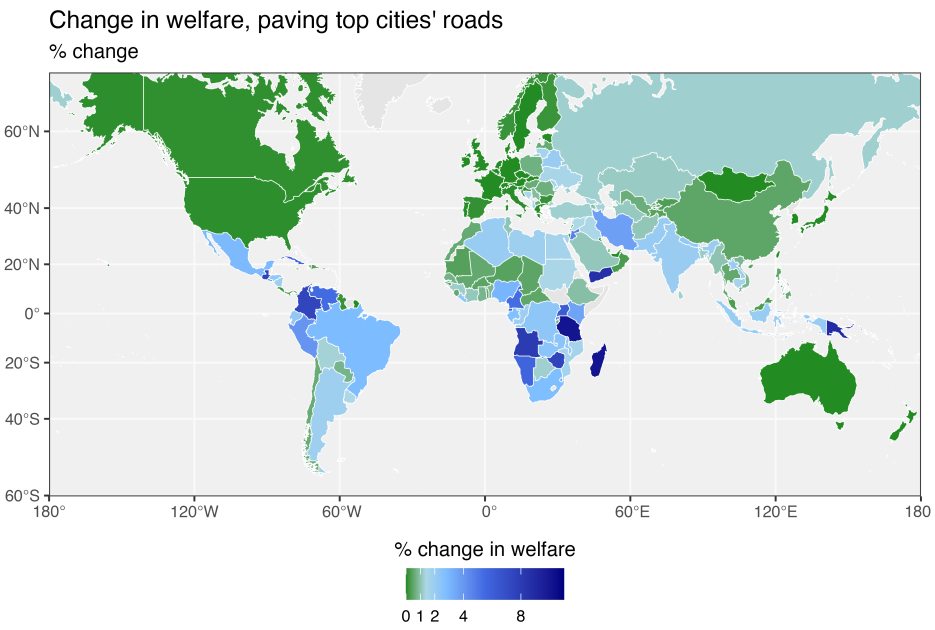


Figure 16: Welfare gain across countries from paving top cities roads, with a budget capped at 1% of GDP.

controls attenuate the coefficient significantly. Conditional on these controls, a 10pp increase in the share of roads paved is associated with a 3.3% decrease in τ_i (p-value $<$0.01).

As small cities are on average the ones with the least paved roads, a policy that aimed at paving the least-paved cities could be potentially misallocative, targeting investment at the least productive cities. To assess the effects of road paving globally, I simulate a policy of targeted investment at the most productive cities in a nation's urban system. To do so, I fix each country's road paving budget at 1% of baseline GDP. I then start by paving roads in the biggest city, and then the second biggest city, and so on, until each city reaches the average level of paved urban roads in the United States (96%). I move down the city size distribution until the paving budget is exhausted or all roads are paved. I assume each new kilometer of paved road costs $227,800 USD, which is the median cost for small projects in African Development Bank (2014). I use this estimate, along with OSRM data on the total length of urban roadways in a city, to compute the costs of the policy.

Figure 16 displays nations welfare changes as a result of the road paving counterfactual. In the rich world, almost all urban roads are paved, and the policy has no effective bite. In the developing world, many countries gain as a result of the policy. However, not all developing countries' workers experience welfare gains in

excess of 1%, making the policy not cost-effective in these countries (on the map, these countries are colored green). These nations are primarily Saharan nations like Chad, Mali, and Niger. These nations have a very low share of urban roads paved at baseline (around 23%) – the issue is they do not have very productive cities, so the general equilibrium gains from the policy are small. On average the share of the population living in cities in these countries rises by only 0.03 percentage points, compared to 0.53pp for the rest of sub-Saharan Africa. In nations like China, the most productive cities are paved, so the policy targets smaller, less productive cities and is not cost-effective. In Appendix Figure A8, I map the population change across cities in sub-Saharan Africa. Targeted cities expand as a result of the policy, while untargeted cities shrink, as their workers leave for better opportunities elsewhere. As some top cities are already well-paved at baseline, in some nations, the policy increases the size of second-tier cities more in percentage terms. For example, in Nigeria, cities like Onitsha or Owerri grow more in percentage terms than Lagos.

6 Conclusion

This paper argues that developing countries have failed to reap the full benefits of urban agglomeration because the costs of urban scale in such nations are too high. Basic facts about cities suggest the fact that the urban costs are high in low-income nations, relative to high-income nations. First, cities in developing countries have accommodated demand for floorspace by building horizontally rather than vertically, suggesting that relative to the rich world, there are technological or regulatory barriers to vertical expansion. Finally, there is more built volume in the urban cores of developing cities relative to their peripheries than there is in rich-world cities. This pattern of development may be driven by either an elastic supply of floorspace or greater demand for cities' downtowns, stemming from poor transportation infrastructure.

To quantify differences in urban technology across space, and how they matter for economic development, I develop a general equilibrium model of cities and an agricultural sector linked through trade and migration. In the model, cities differ in terms of their urban technology: their ability to build vertically, horizontally, and the costs of commuting within them. All three components of a city's urban technology govern the level of urban costs in a city and its urban cost elasticity. The urban cost elasticity is a land-supply elasticity weighted average of congestion forces that operate through the floorspace market and through transportation costs. The

model implies several estimating equations to recover components of the urban technology. My estimates are consistent with the patterns described above: floorspace supply is less elastic in low-income nations, but land is more elastically supplied in low-income nations. However, commuting costs are much higher in low-income nations. Consequently, urban cost elasticities are 35% higher in the low-income nations compared to high-income nations. Using my estimates, I recover building productivity across cities. Building productivity is on average over 3 log points lower in the developing world, compared to the rich world.

I use my quantitative model to explore how these differences in urban costs around the world matter for economic development. Lowering the average level of the urban cost elasticity around the world to the average level in the United States delivers large welfare gains globally. These gains are largest in the developing world, where a third of the gains come from the general equilibrium reallocation of workers out of agriculture towards cities, and the sorting of workers to more productive cities. This counterfactual demonstrates how lower urban costs can foster economic development, and illustrates the role of urban costs play in shaping a nation's aggregate productivity. To understand what available policies can reduce urban costs, I explore the efficacy of urban road paving around the world. In lower middle-income nations, well-targeted road paving is a cost-effective way of reducing urban costs, delivering modest welfare gains in across nations in the developing world.

The framework I have provided could additionally be leveraged to think about several questions in macroeconomics. First, the model could be augmented to include international trade, and used to study how a nation's openness to trade affects the urbanization process. Second, the framework could be used to better understand urbanization via structural transformation by incorporating richer heterogeneity with the agricultural sector, internal bilateral migration costs, and nonhomothetic preferences.

Overall, this study highlights the importance of city-level characteristics in explaining differences in macroeconomic aggregates. Nations with better urban technology have higher aggregate productivity because they can better reap the benefits of cities' productivity advantages. This underscores the importance of research that focuses on how to improve cities around the world. Cities are rapidly changing, but several important features of cities have been absent in this study and demand further attention. First, modes of transit vary across cities, with people in the developing world increasingly using informal transit options and motorcycles. Policy that allows cities to better adapt to these changes in transportation technology may be an effective way

of reducing transportation costs. Second, many developing cities have slums, which governments routinely clear. Whether these slums decrease urban productivity or provide access to cities for people who otherwise would not live in them remains an open question (Monge-Naranjo et al., 2022). As cities are replete with externalities, understanding the costs and benefits of urban policy interventions remains an important agenda. This study has shown that differences in urban costs across cities and nations have large macroeconomic consequences: in short, when it comes to improving the quality of cities, the stakes are large.

References

Ades, Alberto F and Edward L Glaeser (1995). "Trade and circuses: explaining urban giants". In: *The Quarterly Journal of Economics* 110.1, pp. 195–227.

African Development Bank (2014). *Study on Road Infrastructure Costs: Analysis of Unit Costs and Cost Overruns of Road Infrastructure Projects in Africa*. Tech. rep.

Aggarwal, Shilpa (2018). "Do rural roads create pathways out of poverty? Evidence from India". In: *Journal of Development Economics* 133, pp. 375–395.

Ahlfeldt, Gabriel M, Nathaniel Baum-Snow, and Remi Jedwab (2023). "The skyscraper revolution: Global economic development and land savings". In.

Ahlfeldt, Gabriel M and Daniel P McMillen (2014). *New estimates of the elasticity of substitution between land and capital*. Lincoln Institute of Land Policy.

Ahlfeldt, Gabriel M and Elisabetta Pietrostefani (2019). "The economic effects of density: A synthesis". In: *Journal of Urban Economics* 111, pp. 93–107.

Ahlfeldt, Gabriel M, Stephen J Redding, Daniel M Sturm, and Nikolaus Wolf (2015). "The economics of density: Evidence from the Berlin Wall". In: *Econometrica* 83.6, pp. 2127–2189.

Akbar, Prottoy, Victor Couture, Gilles Duranton, and Adam Storeygard (2023a). "Mobility and congestion in urban India". In: *American Economic Review* 113.4, pp. 1083–1111.

— (2023b). *The fast, the slow, and the congested: Urban transportation in rich and poor countries*. Tech. rep. National Bureau of Economic Research.

Albouy, David, Gabriel Ehrlich, and Yingyi Liu (2016). *Housing demand, cost-of-living inequality, and the affordability crisis*. Tech. rep. National Bureau of Economic Research.

Allen, Treb and Costas Arkolakis (2014). "Trade and the Topography of the Spatial Economy". In: *The Quarterly Journal of Economics* 129.3, pp. 1085–1140.

— (2022). "The welfare effects of transportation infrastructure improvements". In: *The Review of Economic Studies* 89.6, pp. 2911–2957.

Allen, Treb, Costas Arkolakis, and Xiangliang Li (forthcoming). "On the Equilibrium Properties of Spatial Models". In: *American Economic Review: Insights*. Forthcoming.

Au, Chun-Chung and J Vernon Henderson (2006). "Are Chinese cities too small?" In: *The Review of Economic Studies* 73.3, pp. 549–576.

Bajzik, Josef, Tomas Havranek, Zuzana Irsova, and Jiri Schwarz (2020). "Estimating the Armington elasticity: The importance of study design and publication bias". In: *Journal of International Economics* 127, p. 103383.

Balboni, Clare, Gharad Bryan, Melanie Morten, and Bilal Siddiqi (2020). "Transportation, gentrification, and urban mobility: The inequality effects of place-based policies". In: *Preliminary Draft* 3.

Baum-Snow, Nathaniel and Lu Han (2024). "The microgeography of housing supply". In: *Journal of Political Economy* 132.6, pp. 1897–1946.

Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud (2014). "Productive cities: Sorting, selection, and agglomeration". In: *Journal of Political Economy* 122.3, pp. 507–553.

Boeing, Geoff (2024). "Modeling and Analyzing Urban Networks and Amenities with OSMnx". In: *Working paper*.

Bordeu, Olivia (2023). "Commuting infrastructure in fragmented cities". In: *Job Market Paper, University of Chicago Booth School of Business* 6.

Brueckner, Jan K (1982). "A note on sufficient conditions for negative exponential population densities". In: *Journal of Regional Science* 22.3, pp. 353–359.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak (2014). "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh". In: *Econometrica* 82.5, pp. 1671–1748.

Bryan, Gharad, Edward Glaeser, and Nick Tsivanidis (2020). "Cities in the developing world". In: *Annual Review of Economics* 12, pp. 273–297.

Bryan, Gharad and Melanie Morten (2019). "The aggregate productivity effects of internal migration: Evidence from Indonesia". In: *Journal of Political Economy* 127.5, pp. 2229–2268.

Cattaneo, Matias D, Richard K Crump, Max H Farrell, and Yingjie Feng (2024). "On binscatter". In: *American Economic Review* 114.5, pp. 1488–1514.

Chari, Amalavoyal, Elaine M Liu, Shing-Yi Wang, and Yongxiang Wang (2021). "Property rights, land misallocation, and agricultural efficiency in China". In: *The Review of Economic Studies* 88.4, pp. 1831–1862.

Chauvin, Juan Pablo, Edward Glaeser, Yueran Ma, and Kristina Tobio (2017). "What is different about urbanization in rich and poor countries? Cities in Brazil, China, India and the United States". In: *Journal of Urban Economics* 98, pp. 17–49.

Coeurdacier, Nicholas, Florian Oswald, and Marc Teignier (2022). "Structural Change, Land Use and Urban Expansion". In.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (2011). "The identification of agglomeration economies". In: *Journal of economic geography* 11.2, pp. 253–266.

— (2019). "The costs of agglomeration: House and land prices in French cities". In: *The Review of Economic Studies* 86.4, pp. 1556–1589.

— (2021). "The production function for housing: Evidence from France". In: *Journal of Political Economy* 129.10, pp. 2766–2816.

Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, and Sébastien Roux (2010). "Estimating agglomeration economies with history, geology, and worker effects". In: *Agglomeration economics*. University of Chicago Press, pp. 15–66.

Combes, Pierre-Philippe and Laurent Gobillon (2015). "The empirics of agglomeration economies". In: *Handbook of regional and urban economics*. Vol. 5. Elsevier, pp. 247–348.

Comin, Diego, Danial Lashkari, and Martí Mestieri (2021). "Structural change with long-run income and price effects". In: *Econometrica* 89.1, pp. 311–374.

Correia, Sergio, Paulo Guimarães, and Tom Zylkin (2020). "Fast Poisson estimation with high-dimensional fixed effects". In: *The Stata Journal* 20.1, pp. 95–115.

Coşar, A Kerem and Pablo D Fajgelbaum (2016). "Internal geography, international trade, and regional specialization". In: *American Economic Journal: Microeconomics* 8.1, pp. 24–56.

Crews, Levi (2023). "A dynamic spatial knowledge economy". In: *Working Paper*.

Davis, Morris A and François Ortalo-Magné (2011). "Household expenditures, wages, rents". In: *Review of Economic Dynamics* 14.2, pp. 248–261.

De La Roca, Jorge and Diego Puga (2017). "Learning by working in big cities". In: *The Review of Economic Studies* 84.1, pp. 106–142.

Desmet, Klaus, Dávid Krisztián Nagy, and Esteban Rossi-Hansberg (2017). "Asia's Geographic Development". In: *Asian Development Review* 34.2, pp. 1–24.

— (2018). "The geography of development". In: *Journal of Political Economy* 126.3, pp. 903–983.

Desmet, Klaus and Esteban Rossi-Hansberg (2013). "Urban accounting and welfare". In: *American Economic Review* 103.6, pp. 2296–2327.

Dingel, Jonathan I (2017). "The determinants of quality specialization". In: *The Review of Economic Studies* 84.4, pp. 1551–1582.

Ducruet, César, Réka Juhász, Dávid Krisztián Nagy, and Claudia Steinwender (2024). "All aboard: The effects of port development". In: *Journal of International Economics* 151, p. 103963.

Duranton, Gilles and Diego Puga (2020). "The economics of urban density". In: *Journal of economic perspectives* 34.3, pp. 3–26.

Epple, Dennis, Brett Gordon, and Holger Sieg (2010). "A new approach to estimating the production function for housing". In: *American Economic Review* 100.3, pp. 905–924.

Esch, Thomas, Elisabeth Brzoska, Stefan Dech, Benjamin Leutner, Daniela Palacios-Lopez, Annekatrin Metz-Marconcini, Mattia Marconcini, Achim Roth, and Julian Zeidler (2022). "World Settlement Footprint 3D-A first three-dimensional survey of the global building stock". In: *Remote sensing of environment* 270, p. 112877.

Esch, Thomas, Julian Zeidler, Daniela Palacios-Lopez, Mattia Marconcini, Achim Roth, Milena Mönks, Benjamin Leutner, Elisabeth Brzoska, Annekatrin Metz-Marconcini, Felix Bachofer, et al. (2020). "Towards a large-scale 3D modeling of the built environment—joint analysis of TanDEM-X, Sentinel-2 and open street map data". In: *Remote Sensing* 12.15, p. 2391.

Finlay, John and Trevor C. Williams (2022). "Housing Demand, Inequality, and Spatial Sorting". In: Working Paper.

Florczyk, A., C. Corbane, M. Schiavina, M. Pesaresi, L. Maffenini, M. Melchiorri, P. Politis, F. Sabo, S. Freire, D. Ehrlich, T. Kemper, P. Tommasi, D. Airaghi, and L. Zanchetta (2019). *GHS Urban Centre Database 2015, multitemporal and multidimensional attributes, R2019A*.

Frolking, Steve, Richa Mahtta, Tom Milliman, Thomas Esch, and Karen C Seto (2024). "Global urban structural growth shows a profound shift from spreading out to building up". In: *Nature Cities*, pp. 1–12.

Gaduh, Arya, Tadeja Gračner, and Alexander D Rothenberg (2022). "Life in the slow lane: Unintended consequences of public transit in Jakarta". In: *Journal of urban economics* 128, p. 103411.

Gibson, John, Susan Olivia, Geua Boe-Gibson, and Chao Li (2021). "Which night lights data should we use in economics, and where?" In: *Journal of Development Economics* 149, p. 102602.

Gollin, Douglas (2014). "The Lewis model: A 60-year retrospective". In: *Journal of Economic Perspectives* 28.3, pp. 71–88.

Gollin, Douglas, Martina Kirchberger, and David Lagakos (2021). "Do urban wage premia reflect lower amenities? Evidence from Africa". In: *Journal of Urban Economics* 121, p. 103301.

Gollin, Douglas, David Lagakos, and Michael E Waugh (2014). "The agricultural productivity gap". In: *The Quarterly Journal of Economics* 129.2, pp. 939–993.

Gomtsyan, David (2024). "Building the city under financial frictions". In: *Review of Economic Dynamics* 52, pp. 70–83.

Gonzalez-Navarro, Marco and Climent Quintana-Domeque (2016). "Paving streets for the poor: Experimental analysis of infrastructure effects". In: *Review of Economics and Statistics* 98.2, pp. 254–267.

Hanson, Gordon H (1998). "Regional adjustment to trade liberalization". In: *Regional science and urban economics* 28.4, pp. 419–444.

Harari, Mariaflavia (2020). "Cities in bad shape: Urban geometry in India". In: *American Economic Review* 110.8, pp. 2377–2421.

Harris, John R and Michael P Todaro (1970). "Migration, unemployment and development: a two-sector analysis". In: *The American economic review* 60.1, pp. 126–142.

Henderson, J Vernon (1974). "The sizes and types of cities". In: *The American Economic Review* 64.4, pp. 640–656.

Henderson, J Vernon, Adam Storeygard, and David N Weil (2012). "Measuring economic growth from outer space". In: *American economic review* 102.2, pp. 994–1028.

Herrendorf, Berthold and Todd Schoellman (2018). "Wages, human capital, and barriers to structural transformation". In: *American Economic Journal: Macroeconomics* 10.2, pp. 1–23.

Hong, Guangbin (2024). "Understanding the Urban Earnings Growth Premium". In.

Jacobs, Jane (1969). *The Economy of Cities*. New York: Random House.

Jedwab, Remi, Prakash Loungani, and Anthony Yezer (2021). "Comparing cities in developed and developing countries: Population, land area, building height and crowding". In: *Regional Science and Urban Economics* 86, p. 103609.

Kreindler, Gabriel, Arya Gaduh, Tilman Graff, Rema Hanna, and Benjamin A Olken (2023). *Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta*. Tech. rep. National Bureau of Economic Research.

Kummu, Matti, Maija Taka, and Joseph HA Guillaume (2018). "Gridded global datasets for gross domestic product and Human Development Index over 1990–2015". In: *Scientific data* 5.1, pp. 1–15.

Lagakos, David (2020). "Urban-rural gaps in the developing world: Does internal migration offer opportunities?" In: *Journal of Economic perspectives* 34.3, pp. 174–192.

Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Corey Vernot, and Michael E Waugh (2020). "Migration costs and observational returns to migration in the developing world". In: *Journal of Monetary Economics* 113, pp. 138–154.

Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E Waugh (2023). "The welfare effects of encouraging rural–urban migration". In: *Econometrica* 91.3, pp. 803–837.

Lagakos, David and Michael E Waugh (2013). "Selection, agriculture, and cross-country productivity differences". In: *American Economic Review* 103.2, pp. 948–980.

Lall, Somik, Mathilde Lebrand, Hogeun Park, Daniel Sturm, and Anthony Venables (2021a). *Pancakes to Pyramids: City form to promote sustainable growth*. World Bank.

Lall, Somik, Mathilde Lebrand, and Maria Edisa Soppelsa (2021b). "The Evolution of City Form". In.

Lewis, W. Arthur (May 1954). "Economic Development with Unlimited Supplies of Labour". In: *The Manchester School* 22.2, pp. 139–191.

Lhuillier, Hugo (2023). "Should I Stay or Should I Grow?" In: *Working Paper*.

Liotta, Charlotte, Vincent Viguié, and Quentin Lepetit (2022). "Testing the monocentric standard urban model in a global sample of cities". In: *Regional Science and Urban Economics* 97, p. 103832.

Lucas, Robert E. (1988). "On the mechanics of economic development". In: *Journal of monetary economics* 22.1, pp. 3–42.

Malpezzi, Stephen (1999). "Economic analysis of housing markets in developing and transition economies". In: *Handbook of regional and urban economics* 3, pp. 1791–1864.

McIntosh, Craig, Tito Alegría, Gerardo Ordóñez, and René Zenteno (2018). "The neighborhood impacts of local infrastructure investment: Evidence from urban Mexico". In: *American Economic Journal: Applied Economics* 10.3, pp. 263–286.

Mills, Edwin S. (1972). *Studies in the Structure of the Urban Economy*. Baltimore, MD: Johns Hopkins Press.

Mills, Edwin S. and Jee Peng Tan (1980). "A comparison of urban population density functions in developed and developing countries". In: *Urban studies* 17.3, pp. 313–321.

Monge-Naranjo, Alexander, Pedro Cavalcanti Ferreira, and Luciene Torres de Mello Pereira (2022). "Of Cities and Slums". In: *Working paper*.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg (Dec. 2018). "Commuting, Migration, and Local Employment Elasticities". In: *American Economic Review* 108.12, pp. 3855–90.

Morris, Carl N (1983). "Parametric empirical Bayes inference: theory and applications". In: *Journal of the American statistical Association* 78.381, pp. 47–55.

Munshi, Kaivan and Mark Rosenzweig (2016). "Networks and misallocation: Insurance, migration, and the rural-urban wage gap". In: *American Economic Review* 106.01, pp. 46–98.

Murray, Michael P and Guoqing Sun (2017). "The demand for space in China". In: *Journal of Urban Economics* 98, pp. 214–222.

Nagy, Dávid Krisztián (2022). "Trade and urbanization: Evidence from Hungary". In: *American Economic Journal: Microeconomics* 14.3, pp. 733–790.

NASA JPL (2020). *NASADEM Merged DEM Global 1 arc second V001*. NASA EOSDIS Land Processes DAAC. DOI: 10.5067/MEaSUREs/NASADEM/NASADEM_HGT.001.

Olken, Benjamin A (2007). "Monitoring corruption: evidence from a field experiment in Indonesia". In: *Journal of political Economy* 115.2, pp. 200–249.

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

Ritchie, Hannah and Max Roser (2019). "Half of the world's habitable land is used for agriculture". In: *Our World in Data*. https://ourworldindata.org/global-land-for-agriculture.

Roberts, Mark, Frederico Gil Sander, and Sailesh Tiwari (2019). *Time to ACT: Realizing Indonesia's Urban Potential*.

Sahai, Harshil and Michael Bailey (2022). "Social Networks and Internal Migration: Evidence from Facebook in India". In: *Working Paper*.

Saiz, Albert (2010). "The geographic determinants of housing supply". In: *The Quarterly Journal of Economics* 125.3, pp. 1253–1296.

Silva, J. M. C. Santos and Silvana Tenreyro (2006). "The Log of Gravity". In: *The Review of Economics and Statistics* 88.4, pp. 641–658.

Storeygard, Adam (2016). "Farther on down the road: transport costs, trade and urban growth in sub-Saharan Africa". In: *The Review of economic studies* 83.3, pp. 1263–1295.

Suárez Serrato, Juan Carlos and Owen Zidar (2016). "Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms". In: *American Economic Review* 106.9, pp. 2582–2624.

Thenkabail, P. S., P. G. Teluguntla, J. Xiong, A. Oliphant, R. G. Congalton, M. Ozdogan, M. K. Gumma, J. C. Tilton, C. Giri, C. Milesi, A. Phalke, R. Massey, K. Yadav, T. Sankey, Y. Zhong, I. Aneece, and D. Foley (2021). *Global cropland-extent product at 30-m resolution (GCEP30) derived from Landsat satellite time-series data for the year 2015 using multiple machine-learning algorithms on Google Earth Engine cloud*. Professional Paper 1868. U.S. Geological Survey, p. 63. DOI: 10.3133/pp1868.

Tsivanidis, Nick (2023). "Evaluating the impact of urban transit infrastructure: Evidence from bogota's transmilenio". In: *Working Paper*.

Wong, Ho Lun, Renfu Luo, Linxiu Zhang, and Scott Rozelle (2013). "Providing quality infrastructure in rural villages: The case of rural roads in China". In: *Journal of Development Economics* 103, pp. 262–274.

Wooldridge, Jeffrey M (2010). *Econometric analysis of cross section and panel data*. MIT press.

Young, Alwyn (2013). "Inequality, the urban-rural gap, and migration". In: *The Quarterly Journal of Economics* 128.4, pp. 1727–1785.

Zanaga, Daniele, Ruben Van De Kerchove, Dirk Daems, Wanda De Keersmaecker, Carsten Brockmann, Grit Kirches, Jan Wevers, Oliver Cartus, Maurizio Santoro, Steffen Fritz, Myroslava Lesiv, Martin Herold, Nandin-Erdene Tsendbazar, Panpan Xu, Fabrizio Ramoino, and Olivier Arino (2022). *ESA WorldCover 10 m 2021 v200*. Version v200. Zenodo. DOI: 10.5281/zenodo.7254221. URL: https://doi.org/10.5281/zenodo.7254221.

A Additional tables and figures

A.1 Additional figures

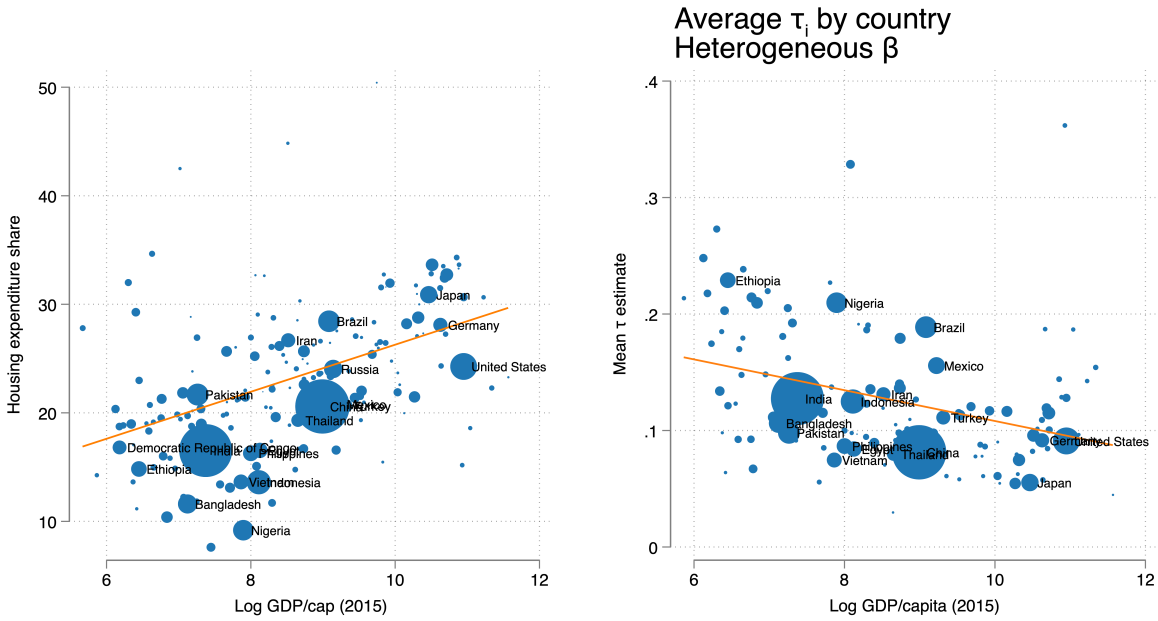


Figure A1: Left: ICP housing expenditure share using data from the World Bank. Right: alternative τ estimates

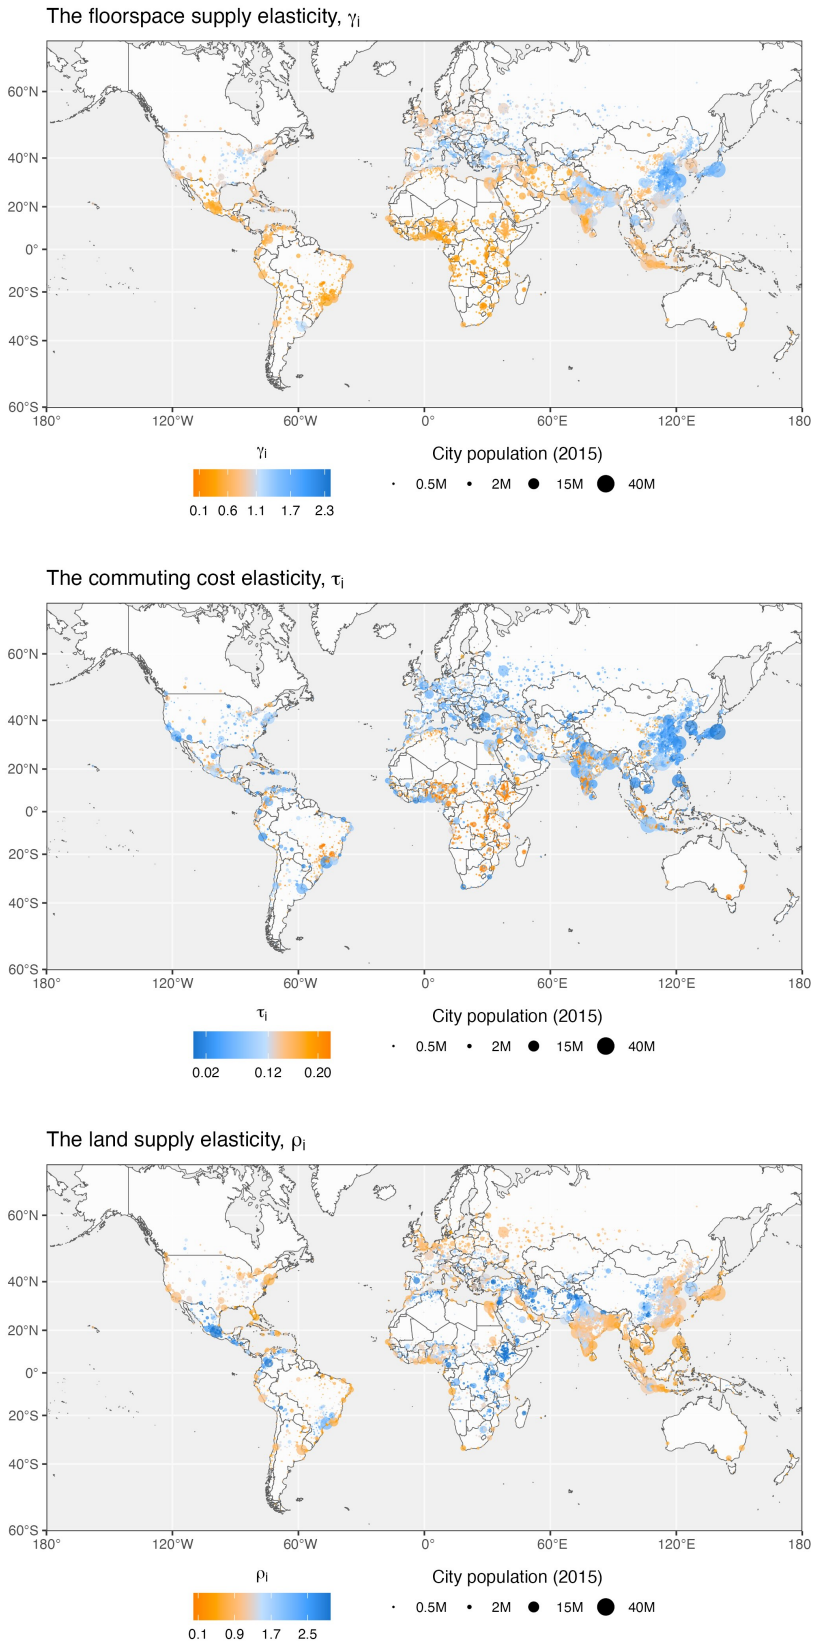


Figure A2: Spatial distribution of floorspace supply elasiticites, γ_i, commuting cost elasticity estimates, τ_i, and land supply elasticities ρ_i. Values for τ_i and ρ_i are right censored around the 98th percentile of the distribution, at values 0.22 and 3, respectively.

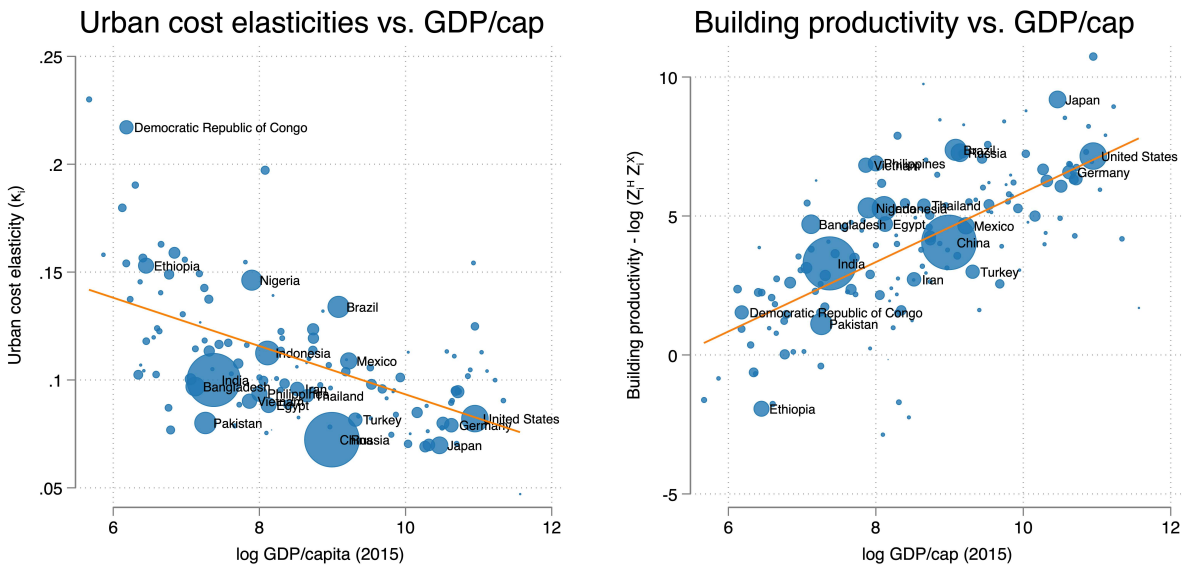


Figure A3: Left: The average urban cost elasticity κ_i across cities in a nation versus a nation's log GDP/capita. Right: average log building productivity ($Z_i^H Z_i^X$) versus a nation's log GDP/capita.

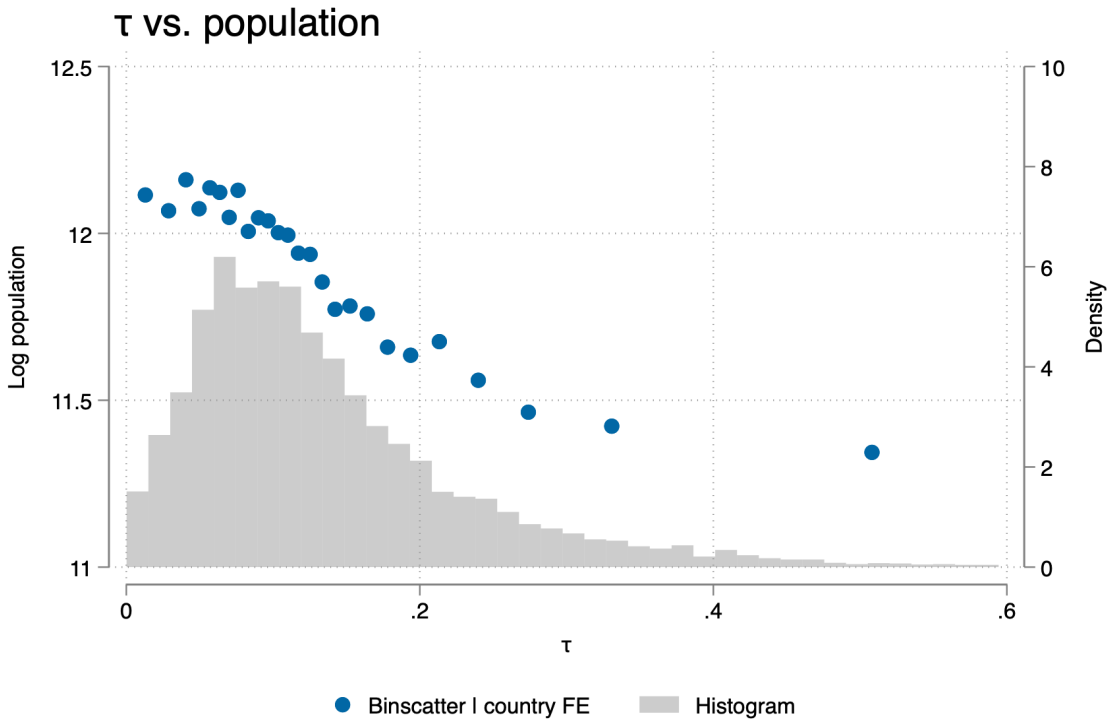


Figure A4: scatter of τ_i vs log population net country fixed effects (using the methodology of Cattaneo et al. (2024). Histogram of τ_i shown in grey.

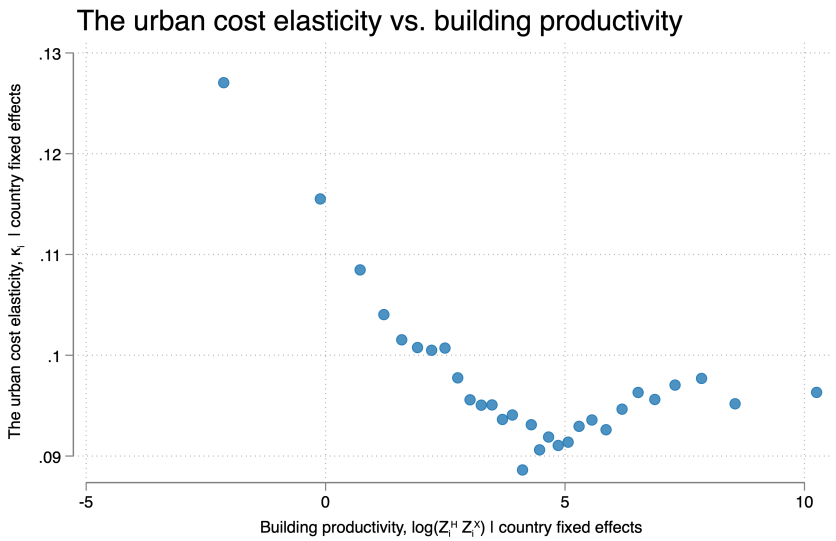


Figure A5: Conditional binscatter of the urban cost elasticity, κ_i versus building productivity, $Z_i^H Z_i^X$, conditional on country fixed effects.

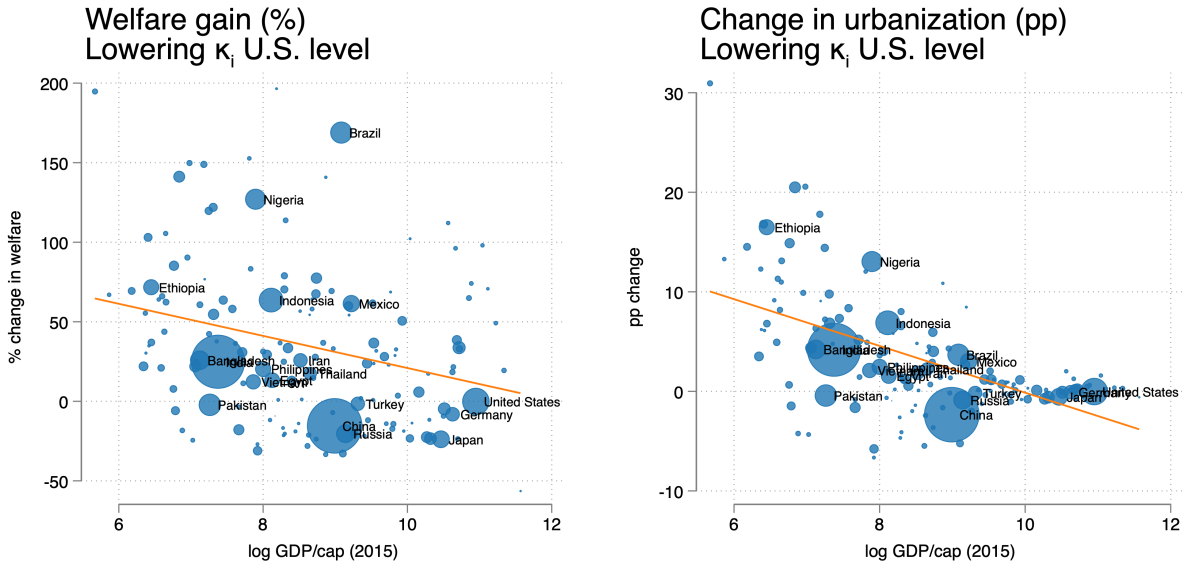


Figure A6: Welfare and urbanization effects from changing the urban cost elasticity, κ_i to the U.S. level.

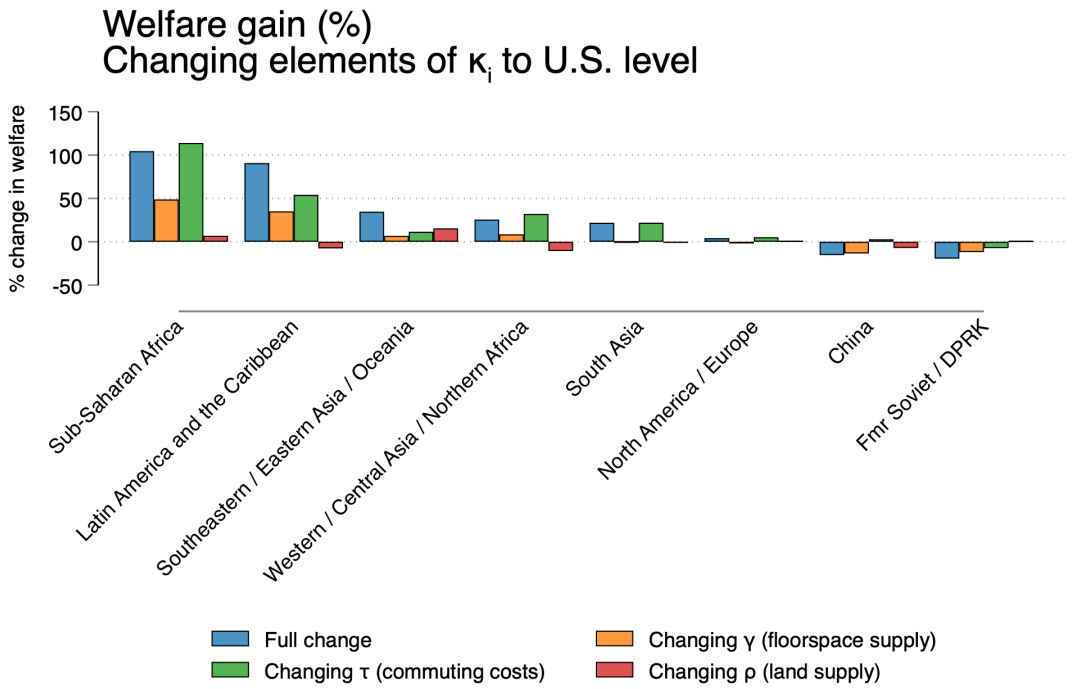


Figure A7: Welfare effects across regions from changing different components of the urban cost elasticity.

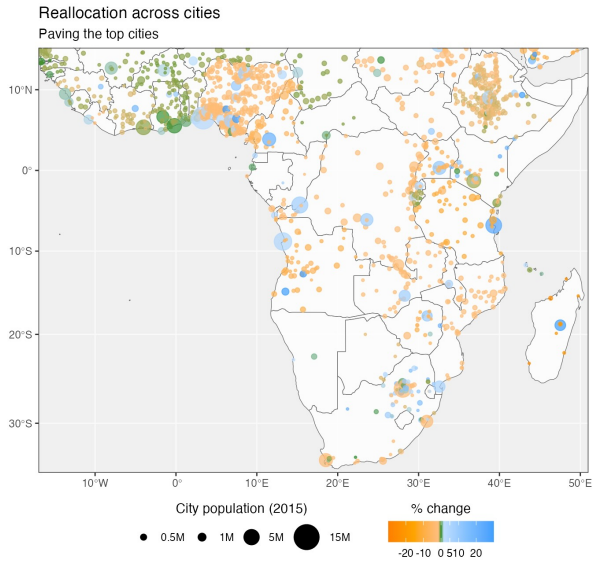


Figure A8: Population change in sub-Saharan African cities from paving top cities' roads.

A.2 Additional tables

	$\ln$(Volume/Person)	$\ln$(Height)	$\ln$(Dev. Area/Person)	$\ln$(Height) $-$ $\ln$(Dev. Area/Person)
Raw	0.334	0.225	0.110	0.115
	(0.027)	(0.019)	(0.025)	(0.034)
	[0.412]	[0.451]	[0.085]	[0.056]
Pop-adjusted	0.324	0.218	0.106	0.111
	(0.025)	(0.018)	(0.022)	(0.032)
	[0.449]	[0.436]	[0.096]	[0.057]
Cities	15,110 (pop adjustment)			
Countries	188	188	188	188

Table A1: Robust (HC1) standard errors in parentheses. R^2 in brackets. Pop adjustment: residualized on second order polynomial in log population and including a dummy for primacy, done at city level. Column 4 tests whether height and area respond equally to income.

Region	κ_i	τ_i	γ_i	ρ_i
China	0.071	0.084	1.346	1.402
Former Soviet / DPRK	0.071	0.070	1.272	1.206
South Asia	0.097	0.149	1.060	1.317
Latin America and the Caribbean	0.115	0.161	0.620	1.505
North America and Europe	0.082	0.086	1.062	1.136
Southeastern/Eastern Asia and Oceania	0.097	0.107	0.946	0.872
Sub-Saharan Africa	0.149	0.254	0.454	2.087
Western/Central Asia and Northern Africa	0.092	0.123	0.879	1.849

Table A2: Averages of the urban cost elasticity, κ_i and its components, the commuting cost elasticity (τ_i), the floorspace supply elasticity (γ_i) and the land supply elasticity (ρ_i) by region.

	τ_i	γ_i	ρ_i
Coefficient	0.321	-0.034	-0.003
	(0.003)	(0.001)	(0.000)
Partial R-squared	0.473	0.242	0.022

Table A3: Partial R-squareds from a regression of κ_i on its components. Each column reports a regression coefficient and partial R-squared value for the different elements of κ_i from a multivariate regression of κ_i on τ_i, γ_i and ρ_i. Conventional standard errors in parentheses.

	Speed near city center				Speed indices			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log dowtown speed (midnight)	-1.064	-0.374	-1.899	-1.222				
	(0.277)	(0.340)	(0.391)	(0.476)				
log dowtown speed (midday)	0.917	0.239	1.478	0.660				
	(0.238)	(0.307)	(0.278)	(0.417)				
Uncongested speed index					-1.610	-0.219	-3.103	-1.726
					(0.459)	(0.524)	(0.769)	(0.906)
Speed index					1.679	0.185	2.919	1.210
					(0.463)	(0.535)	(0.623)	(0.811)
log pop		-0.090		-0.053		-0.099		-0.067
		(0.038)		(0.044)		(0.036)		(0.043)
log population/km2		0.002		-0.162		0.016		-0.162
		(0.052)		(0.085)		(0.051)		(0.087)
1(primate city)		-0.152		-0.161		-0.155		-0.160
		(0.117)		(0.159)		(0.120)		(0.163)
Observations	856	856	856	856	856	856	856	856
R-squared	0.02	0.04	0.26	0.27	0.02	0.04	0.25	0.27
Country FE			✓	✓			✓	✓

Table A4: Correlation with Akbar et al. (2023b) city speed variables. Dependent variable: $\log \tau_i$. Robust standard errors in parentheses.

	κ_i		$\log \kappa$	
	(1)	(2)	(3)	(4)
log GDP/cap (country)	-0.013		-0.081	
	(0.002)		(0.013)	
log GDP/cap (city)		-0.006		-0.036
		(0.001)		(0.003)
log population (city)		-0.006		-0.036
		(0.000)		(0.003)
Observations	127	9,358	127	9,358
R-squared	0.25	0.39	0.25	0.39
R-squared (within)		0.04		0.03
Country FE		✓		✓

Table A5: Income and population correlations with κ_i. Robust standard errors in parentheses.

	(1) Shipment value
Log dist	-0.923
	(0.022)
N	4,817

Table A6: Gravity regression of the value of intercity road shipments against distance, conditional on origin and destination fixed effects in the 2017 U.S. Commodity Flows Survey. Robust standard errors in parentheses. Distance is defined as $1 + d_{ij}$ where d_{ij} is measured in kilometers.

	(1)	(2)	(3)
Share of roads paved	-0.790	-0.709	-0.328
	(0.033)	(0.071)	(0.117)
Log total street length		-0.255	-0.142
		(0.064)	(0.079)
Log average number of lanes		-0.094	0.021
		(0.098)	(0.112)
Log number of intersections		0.106	0.011
		(0.055)	(0.074)
Log average road width		-0.015	-0.015
		(0.025)	(0.026)
1(primate)		-0.214	-0.217
		(0.072)	(0.082)
Log population		0.015	-0.001
		(0.023)	(0.026)
Observations	5,641	1,604	1,581
R-squared	0.09	0.14	0.29
R-squared (within)			0.08
Country FE			✓

Table A7: Dependent variable: $\log \tau_i$. Robust standard errors in parentheses.

B Full model derivation

Environment

The economy of nation n is composed of a mass of L_n households, and I_n cities indexed by i which produce traded urban varieties, and an agricultural sector denoted by a, which produces a freely traded good which is numeraire.

Cities are monocentric and circular with radius X_i. Each city contains a citywide developer that constructs both urban land and floorspace atop it using the agricultural good and land.

Ex-ante identical households whether to work in the agricultural sector or urban sector, and conditional on working in the urban sector, decide in which city to live and the location x within the city to live. Their location choice depends on a realization of an idiosyncratic preference shock which is iid across cities and the agricultural sector. Households are freely mobile across locations with in each city. Households inelastically supply labor to the sector in which they work and earn a wage w_i, which finances their consumption of urban varieties, the agricultural good, and floorspace. Urban development is absent in the rural sector, and consequently their demand for floorspace is met by directly consuming land which they are also assumed to own, – i.e., they 'live on the farm.'

Preferences

Households have Cobb-Douglas preferences over urban varieties (traded), agriculture (freely traded and the numeraire), and floorspace. For agent ν, their problem is,

$$\max_{i,x,c_u,c_a,h} A_i(x)\left(\frac{c_u}{\alpha}\right)^{\alpha}\left(\frac{\psi^H h}{\beta}\right)^{\beta}\left(\frac{c_a}{1-\alpha-\beta}\right)^{1-\alpha-\beta}\epsilon_i^{\nu} \quad \text{s.t.} \quad \sum_j p_{ji}c_j + q_i(x)h + c^a \leq w_i$$

where,

$$c_u = \left(\sum_{j=1}^{I_n}(c_j)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$$

and urban varieties each have price p_{ji}.

The term ϵ_i^{ν} is an idiosyncratic preference shock which is Fréchet distributed with scale parameter ε iid over locations, where $i = a$ for the agricultural sector and otherwise $i = 1, ..., I_n$ enumerates cities.

Consequently, in any city i and location x, indirect utility is,

$$v_i(x) = A_i(x)\frac{w_i}{P_i^{\alpha}q_i(x)^{\beta}}(\psi^H)^{\beta}$$

where,

$$P_i = \left(\sum_j (p_{ji})^{1-\sigma}\right)^{\frac{1}{1-\sigma}}.$$

Locations in cities are exogenously characterized by their amenity supply functions,

$$A_i(x) = \underbrace{\tilde{A}_i}_{\text{citywide amenities}} \cdot \underbrace{(x)^{-\tau_i}}_{\text{commuting costs}} .$$

where,

$$\tilde{A}_i = A_i \left(1 - \frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)^{-\left(\frac{\beta}{1+\gamma_i}\frac{1}{1+\rho_i} + \frac{\rho_i}{1+\rho_i}\frac{\tau_i}{2}\right)} .$$

This functional form is chosen so that the elasticities γ_i, τ_i do not affect the *level* of utility in any given city.

Citybuilding technology

In each city, there is a continuum of identical developers that can build floorspace per unit of land at each x with the technology,

$$H_i(x) = \left(\frac{\gamma_i}{1+\gamma_i}\right)^{-1} \left(\frac{Z_i^H}{\psi^H}\right)^{\frac{1}{1+\gamma_i}} c_a^{\frac{\gamma_i}{1+\gamma_i}} T_i(x)^{\frac{1}{1+\gamma_i}} .$$

To begin construction, they must pay a fixed cost,

$$F_i(x) = \left(\frac{\psi^X}{Z_i^X \left(\frac{1}{1+\gamma_i}\left(1 - \frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)\right)^{\frac{-\rho_i}{1+\rho_i}}}\right)^{1+1/\rho_i} (x)^{2/\rho_i} .$$

What matters in this functional form is that fixed costs rise in x at rate $2/\rho_i$. The rescaling of the level of fixed costs ensures that the intercept of the land supply curve does not depend on elasticities and is a parameter I can hold fixed.

Subsequently, for a developer who enters at x, their problem is,

$$\max_{c_a} q_i(x)H_i(x) - c^a - F_i(x), \quad \text{such that} \quad H_i(x) = \left(\frac{\gamma_i}{1+\gamma_i}\right)^{-1} \frac{Z_i^H}{\psi^H} c_a^{\frac{\gamma_i}{1+\gamma_i}} T_i(x)^{\frac{1}{1+\gamma_i}}$$

Developers do not enter at x if land rents $r_i(x) < F_i(x)$, where land rents are,

$$r_i(x) = \max_{c^a} q_i(x)H_i(x) - c^a.$$

City partial equilibrium

A city (partial) equilibrium for city i is, given $\{w_i, L_i\}$ and $\{p_{ji}\}$ as well as $A_i(x), Z_i^H, Z_i^X, \psi^H, \psi^X$, and parameters $\{\alpha, \beta, \gamma_i, \tau_i, \rho_i, \sigma\}$, is a set of floorspace prices and quantities, $\{q_i(x), H_i(x)\}$, an urban radius X_i, a population distribution across locations in the city $\{L_i(x)\}$, and a common level of (indirect) utility v_i such that,

1. Freely mobile households maximize, taking wages, prices, and the rebate as given, so that floorspace demand per-person is,

$$h_i(x) = \frac{\beta w_i}{q_i(x)}$$

and a spatial equilibrium holds so that indirect utility everywhere is equal to some citywide value v_i,

$$v_i(x) = v_i \; \forall x \in (0, X_i];$$

2. each developer profit maximizes, optimally choosing $H_i(x)$ given $q_i(x)$;
3. developers freely enter until $r_i(X_i) = F_i(X_i)$;
4. each household is housed somewhere, so that,

$$2\pi \int_0^{X_i} x L_i(x) dx = L_i;$$

5. and the floorspace market clears so that,

$$H_i(x) = h_i(x) L_i(x) \; \forall x \in (0, X_i].$$

Solution to the city equilibrium Starting with household indirect utility,

$$v_i(x) = A_i(x) \frac{w_i}{P_i^{\alpha} q_i(x)^{\beta}} (\psi^H)^{\beta}$$

by spatial equilibrium ($v_i(x) = v_i \; \forall x \in (0, X_i]$), this means,

$$q_i(x) = \left(\frac{A_i(x) w_i P_i^{-\alpha}}{v_i} \right)^{1/\beta}$$

where v_i is the common level of utility achieved at all locations x in the city.

Inverting the first order condition on $H_i(x)$ gives the floorspace supply function,

$$q_i(x) \left(\frac{Z_i^H}{\psi^H} \right)^{\frac{-1}{1+\gamma_i}} (c_a)^{\frac{-1}{1+\gamma_i}} = 1$$
$$\implies H_i(x) = \frac{Z_i^H}{\psi^H} q_i(x)^{\gamma_i}$$

then by housing market clearing,

$$q_i(x) H_i(x) = \beta w_i L_i(x).$$

Plugging in for $H_i(x)$,

$$q_i(x)^{1+\gamma_i} \frac{Z_i^H}{\psi^H} = \beta w_i L_i(x)$$

or,

$$q_i(x) = \left(\frac{\beta w_i L_i(x)}{\frac{Z_i^H}{\psi^H}} \right)^{\frac{1}{1+\gamma_i}}$$

combining with the spatial equilibrium condition gives,

$$\left(\frac{A_i(x)w_iP_i^{-\alpha}}{v_i}\right)^{1/\beta} = \left(\frac{\beta w_i L_i(x)}{\frac{Z_i^H}{\psi^H}}\right)^{\frac{1}{1+\gamma_i}}$$

we can now use this to solve for v_i and consequently $L_i(x)$ using,

$$2\pi \int_0^{X_i} xL_i(x)dx = L_i$$

Inverting the housing market clearing condition with plugged-in supply functions and price gradients to solve for $L_i(x)$ gives,

$$L_i(x) = \frac{A_i(x)^{\frac{1+\gamma_i}{\beta}}}{2\pi \int_0^{X_i} sA_i(s)^{\frac{1+\gamma_i}{\beta}} ds} L_i$$

Consider the functional form on $A_i(x)$,

$$A_i(x) = \tilde{A}_i \cdot x^{-\tau_i}$$

Now consider,

$$\begin{aligned}
L_i(x) &= \frac{x^{-\tau_i\frac{1+\gamma_i}{\beta}}}{2\pi \int_0^{X_i} s^{1-\tau_i\frac{1+\gamma_i}{\beta}} dsl} L_i \\
&= \frac{x^{-\tau_i\frac{1+\gamma_i}{\beta}}}{2\pi \times \left[\frac{1}{2-\tau_i\frac{1+\gamma_i}{\beta}} s^{2-\tau_i\frac{1+\gamma_i}{\beta}}\right]_0^{X_i}} L_i \\
&= \frac{x^{-\tau_i\frac{1+\gamma_i}{\beta}}}{2\pi X_i \frac{1}{2-\tau_i\frac{1+\gamma_i}{\beta}} X_i^{-\tau_i\frac{1+\gamma_i}{\beta}}} L_i \\
&= \left(1 - \frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)\left(\frac{x}{X_i}\right)^{-\tau_i\frac{1+\gamma_i}{\beta}} \frac{L_i}{\pi X_i^2}
\end{aligned}$$

this is the population distribution along any ray from the center of the city ($x = 0$).

Finally, we must solve for X_i. Land rents are a constant share of revenue at the periphery,

$$\frac{\beta}{1+\gamma_i} w_i L_i(X_i) = F_i(X_i)$$

Plugging in for $L_i(X_i)$ and $F_i(X_i)$

$$\frac{\beta}{1+\gamma_i} w_i \left(1 - \frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right) \frac{L_i}{\pi X_i^2} = \left(\frac{\psi^X}{Z_i^X \left(\frac{1}{1+\gamma_i}\left(1 - \frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)\right)^{\frac{-\rho_i}{1+\rho_i}}}\right)^{1+1/\rho_i} (x)^{2/\rho_i}.$$

Solving for area gives the land supply curve,

$$\pi X_i^2 = \pi \left(\frac{Z_i^X}{\psi^X}\right) \left(\frac{\beta}{\pi}\beta w_i L_i\right)^{\frac{\rho_i}{1+\rho_i}}.$$

Characterizing the equilibrium: city level aggregates

Because of the spatial equilibrium condition, it is sufficient to characterize utility in the city with utility at the fringe, $v_i = v_i(X_i)$. So,

$$\begin{aligned}
v_i &= v_i(X_i) \\
&= A_i \cdot X_i^{-\tau} \frac{w_i}{P_i^\alpha q_i(X_i)^\beta} \\
&= A_i \cdot X_i^{-\tau} \frac{w_i}{P_i^\alpha} \left(\frac{\beta w_i L_i(X_i)}{Z_i^H/\psi^H}\right)^{\frac{-\beta}{1+\gamma_i}} \\
&= A_i \cdot X_i^{-\tau} \frac{w_i}{P_i^\alpha} \left(\frac{\beta w_i}{Z_i^H/\psi^H}\left(1 - \frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)\frac{L_i}{\pi X_i^2}\right)^{\frac{-\beta}{1+\gamma_i}} \\
&= A_i \frac{w_i}{P_i^\alpha} (Z_i^H/\psi^H)^{\frac{\beta}{1+\gamma_i}} \left(1 - \frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)^{\frac{-\beta}{1+\gamma_i}} (\beta/\pi)^{\frac{-\beta}{1+\gamma_i}} (w_i L_i)^{\frac{-\beta}{1+\gamma_i}} (X_i^2)^{2\left(\frac{\beta}{1+\gamma_i} - \frac{\tau_i}{2}\right)}
\end{aligned}$$

Now using,

$$X_i^2 = \left(\frac{Z_i^X}{\psi^X}\right) (w_i L_i)^{\frac{\rho_i}{1+\rho_i}}.$$

and,

$$\tilde{A}_i = A_i \left(1 - \frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)^{-\left(\frac{\beta}{1+\gamma_i}\frac{1}{1+\rho_i} + \frac{\rho_i}{1+\rho_i}\frac{\tau_i}{2}\right)}$$

$$v_i = v_i(X_i) = A_i \left(\frac{Z_i^H}{\psi^H}\right)^{\frac{\beta}{1+\gamma_i}} \left(\frac{Z_i^X}{\psi^X}\right)^{\left(\frac{\beta}{1+\gamma_i} - \frac{\tau_i}{2}\right)} \frac{w_i}{P_i^\alpha} \left(\frac{\beta}{\pi} w_i L_i\right)^{-\left(\frac{\beta}{1+\gamma_i}\frac{1}{1+\rho_i} + \frac{\rho_i}{1+\rho_i}\frac{\tau_i}{2}\right)}$$

How much floorspace volume is built in aggregate? Call aggregate floorspace H_i,

$$\begin{aligned}
H_i &= 2\pi \int_0^{X_i} x H_i(x) dx \\
&= 2\pi \int_0^{X_i} x \frac{Z_i^H}{\psi^H} q_i(x)^{\gamma_i} dx \\
&= 2\pi \frac{Z_i^H}{\psi^H} \int_0^{X_i} x \left(\frac{\beta w_i L_i(x)}{\frac{Z_i^H}{\psi^H}} \right)^{\frac{\gamma_i}{1+\gamma_i}} dx \\
&= 2\pi \left(\frac{Z_i^H}{\psi^H} \right)^{\frac{1}{1+\gamma_i}} \left(\frac{\beta w_i L_i}{\pi X_i^2} \right)^{\frac{\gamma_i}{1+\gamma_i}} (X_i)^{\tau_i \frac{\gamma_i}{\beta}} \left(1 - \frac{\tau_i}{2} \frac{1+\gamma_i}{\beta} \right)^{\frac{\gamma_i}{1+\gamma_i}} \int_0^{X_i} (x)^{1-\tau_i \frac{\gamma_i}{\beta}} dx \\
&= 2\pi \left(\frac{Z_i^H}{\psi^H} \right)^{\frac{1}{1+\gamma_i}} \left(\frac{\beta w_i L_i}{\pi X_i^2} \right)^{\frac{\gamma_i}{1+\gamma_i}} (X_i)^{\tau_i \frac{\gamma_i}{\beta}} \left(1 - \frac{\tau_i}{2} \frac{1+\gamma_i}{\beta} \right)^{\frac{\gamma_i}{1+\gamma_i}} \left[\frac{1}{2 - \tau_i \frac{\gamma_i}{\beta}} X_i^{2-\tau_i \frac{\gamma_i}{\beta}} \right]
\end{aligned}$$

Then,

$$H_i = 2\pi \left(\frac{Z_i^H}{\psi^H} \right)^{\frac{1}{1+\gamma_i}} \left(\frac{\beta w_i L_i}{\pi X_i^2} \right)^{\frac{\gamma_i}{1+\gamma_i}} (X_i^2) \frac{\left(1 - \frac{\tau_i}{2} \frac{1+\gamma_i}{\beta} \right)^{\frac{\gamma_i}{1+\gamma_i}}}{2 - \tau_i \frac{\gamma_i}{\beta}}$$

Which can be rearranged,

$$\frac{H_i}{\pi X_i^2} = \left(\frac{\beta w_i L_i}{\pi X_i^2} \right)^{\frac{\gamma_i}{1+\gamma_i}} \left(\frac{Z_i^H}{\psi^H} \right)^{\frac{1}{1+\gamma_i}} \frac{\left(1 - \frac{\tau_i}{2} \frac{1+\gamma_i}{\beta} \right)^{\frac{\gamma_i}{1+\gamma_i}}}{1 - \frac{\tau_i}{2} \frac{\gamma_i}{\beta}}$$

This implies the relationship between total built volume, GDP, and area,

$$\log H_i = \frac{\gamma_i}{1+\gamma_i} \log w_i L_i + \frac{1}{1+\gamma_i} \log X_i^2 + a_i$$

where $w_i L_i$ is city GDP and ξ_i subsumes city-specific constants.

The total area of the city implies the estimating equation

$$\log \pi X_i^2 = \frac{\rho_i}{1+\rho_i} \log w_i L_i + b_i$$

where b_i subsumes the TFP constant.

Production in urban and rural sectors

I have taken w_i and L_i to so far be exogenous and have not discussed the agricultural sector. I will now characterize both of these objects.

Production in cities is constant returns with external scale economies which capture urban agglomeration

forces,

$$y_i = Z_i^y L_i, \quad Z_i^y = \bar{Z}_i^y \cdot \left(\frac{L_i}{\pi X_i^2}\right)^{\zeta}, \zeta > 0$$

The rural sector combines land and labor to produce agricultural goods,

$$y_a = Z_a (L_a)^{1-\mu} (T_a^y)^{\mu}$$

where T_a^y is the choice of land used for farming.

General equilibrium

A general equilibrium of the model is, given $\{Z_i^y\}, Z_a, T_a$ and parameters $\{\mu, \lambda, \varepsilon\}$ and all primitives of the city-partial equilibrium, a set of wages and a population distribution $\{w_i, L_i\}, w_a, L_a$ such that,

1. Taking wages and floorspace prices as given, households optimizally choose their location (across cities), and floorspace is demanded optimally in the agricultural sector;
2. taking wages and land prices as given, urban and rural producers profit-maximize;
3. the aggregate labor market clears so that,
$$\sum_{i=a,1,\ldots,I_n} L_i = L_n$$
4. the agricultural land market clears, so that land for consumption and production exhausts the fixed supply of agricultural land,
$$T_a^y + h_a L_a = T_a;$$
5. land rents in the agricultural sector are rebated back to agricultural households;
6. the urban goods market clears, so that,
$$y_i = \sum_j c_{ij}$$
where c_{ij} is the total amount of city i's good is demanded by city j;
7. the agricultural goods market clears;
8. and all optimization and market clearing conditions of the city partial equilibrium hold.

General equilibrium solution

First, by properties of the Fréchet distribution, the population distribution across cities is given by,

$$L_i = \frac{(v_i)^{\varepsilon}}{(v_a)^{\varepsilon} + \sum_{j=1}^{I_n} (v_j)^{\varepsilon}} L_n, \quad L_a \propto \frac{(v_a)^{\varepsilon}}{(v_a)^{\varepsilon} + \sum_{j=1}^{I_n} (v_j)^{\varepsilon}} L_n.$$

Both urban and agricultural labor markets are competitive, so profit maximization ensures that wages are marginal products,

$$w_i = p_i Z_i^y, \quad w_a = Z_a (L_a / T_a^y)^{-\mu}$$

Households earn a fraction $(1-\mu)$ of agricultural income Y_a as wage income, while the fraction μ is spent on land for productive use. Total land rents in the agricultural sector, $q_a T_a$ are rebated back to households, so total expenditure on land for floorspace is in equilibrium given by land market clearing,

$$T_a^h + T_a^h = T_a \implies \beta(1-\mu)Y_a + \mu Y_a + \beta q_a T_a = q_a T_a$$

This implies,

$$\left(\frac{\beta(1-\mu)+\mu}{1-\beta}\right)\frac{Y_a}{T_a} = q_a$$

Since productive land demand satisfies,

$$q_a T_a^y = \mu Y_a \implies \frac{\beta(1-\mu)+\mu}{1-\beta}\frac{Y_a}{T_a}T_a^y = \mu Y_a$$

or, that productive land use is a constant fraction of total land use,

$$T_a^y = \frac{\mu(1-\beta)}{\beta(1-\mu)+\mu}T_a = \frac{\mu(1-\beta)}{\mu(1-\beta)+\beta}T_a$$

Therefore, $T_a^h = \frac{\beta}{\mu(1-\beta)+\beta}T_a$.

Wages $\tilde{w}_a$ are marginal products, so,

$$\tilde{w}_a = (1-\mu)Z_a(L_a)^{-\mu}(T_a^y)^{\mu} = (1-\mu)\underbrace{Z_a\left(\frac{\mu(1-\beta)}{\mu(1-\beta)+\beta}T_a\right)^{\mu}}_{\equiv \tilde{Z}_a}(L_a)^{-\mu}$$

Total agricultural compensation per worker, $\tilde{w}_a$, which accounts for the rebate and is measured in units of Y_a is also,

$$\begin{aligned}
w_a &= (1-\mu)\frac{Y_a}{L_a} + \frac{q_a T_a}{L_a} \\
&= \frac{1}{L_a}\left((1-\mu)\frac{Y_a}{L_a} + \frac{q_a T_a}{L_a}\right) \\
&= \frac{1}{L_a}\left((1-\mu)Y_a + q_a T_a\right) \\
&= \frac{1}{L_a}\left((1-\mu)Y_a + \frac{\beta(1-\mu)+\mu}{1-\beta}Y_a\right) \\
&= \frac{Y_a}{L_a}\left(\frac{(1-\mu)(1-\beta)+\beta(1-\mu)+\mu}{1-\beta}\right) \\
&= \frac{Y_a}{L_a}\left(\frac{1}{1-\beta}\right)
\end{aligned}$$

Now to derive indirect utility in the agricultural sector,

$$\begin{aligned}
v_a &= A_a \frac{w_a}{P_a^\alpha q_a^\beta} \\
&= A_a \frac{\frac{Y_a}{L_a}\left(\frac{1}{1-\beta}\right)}{P_a^\alpha \left(\frac{\beta(1-\mu)+\mu}{1-\beta}\frac{Y_a}{T_a}\right)^\beta} \\
&= A_a (T_a)^\beta \frac{\frac{Y_a}{L_a}\left(\frac{1}{1-\beta}\right)}{P_a^\alpha \left(\frac{\beta(1-\mu)+\mu}{1-\beta}\frac{Y_a L_a}{L_a}\right)^\beta} \\
&= A_a (T_a)^\beta \left(\frac{1}{1-\beta}\right)\left(\frac{\beta(1-\mu)+\mu}{1-\beta}\right)^{-\beta} \frac{\frac{Y_a}{L_a}}{P_a^\alpha \left(\frac{Y_a L_a}{L_a}\right)^\beta} \\
&= A_a (T_a)^\beta \left(\frac{1}{1-\beta}\right)\left(\frac{\beta(1-\mu)+\mu}{1-\beta}\right)^{-\beta} \frac{\left(\frac{Y_a}{L_a}\right)^{1-\beta}}{P_a^\alpha} L_a^{-\beta}
\end{aligned}$$

Now, plugging in for the production function, $Y_a/L_a = \tilde{Z}_a (L_a)^{-\mu}$, we have that,

$$v_a = A_a (T_a)^\beta \left(\frac{1}{1-\beta}\right)\left(\frac{\beta(1-\mu)+\mu}{1-\beta}\right)^{-\beta} \frac{\left(\tilde{Z}_a\right)^{1-\beta} (L_a)^{-(\mu(1-\beta)+\beta)}}{P_a^\alpha}$$

which captures the two congestion forces of population in the agricultural sector: bidding up housing prices, and bidding down wages.

Finally, we need to pin down wages in the urban sector through goods market clearing. Agricultural goods market clearing will mechanically be satisfied by Walras' law.

Goods market clearing requires that the value produced in the urban sector, measured in wages, equals economywide expenditure on that good.

For city i,

$$w_i L_i = \sum_{j=1}^{I_n} p_{ij} c_{ij} + p_{ia} c_{ia}$$

where p_{ij} denotes the price paid for one unit of i's good in location j. Assuming trade costs are iceberg τ_{ij} and the output market is competitive, $p_{ij} = \tau_{ij} p_i$. The term c_{ij} denotes j's demand for i's good. As the demand system is CES,

$$c_{ij} = \left(\frac{\tau_{ij} p_i}{P_j}\right)^{-\sigma} \cdot \frac{\alpha w_j L_j}{P_j}$$

Therefore,

$$w_i L_i = \sum_{j=1}^{I_n} \left(\frac{\tau_{ij} p_i}{P_j}\right)^{1-\sigma} \alpha w_j L_j + \left(\frac{\tau_{ia} p_i}{P_a}\right)^{1-\sigma} \frac{\alpha}{1-\beta} w_a L_a.$$

As the labor market is competitive, $p_i = \frac{w_i}{Z_i^y}$, so,

$$w_i L_i = \sum_{j=1}^{I_n} \left(\frac{\tau_{ij} \frac{w_i}{Z_i^y}}{P_j} \right)^{1-\sigma} \alpha w_j L_j + \left(\frac{\tau_{ia} p_i}{P_a} \right)^{1-\sigma} \frac{\alpha}{1-\beta} w_a L_a,$$

and,

$$P_j = \left(\sum_i \left(\tau_{ij} \frac{w_i}{Z_i^y} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}},$$

closing the model.

B.1 Allen et al (2024) proof application

The equilibrium can be characterized as finding the solution to a nonlinear system in $\{w_i, L_i\}$, since X_i, H_i and the population distribution within cities can all be solved in closed form, and prices are marginal costs.

The system resembles that of Allen and Arkolakis (2014) with three differences,

1. Agglomeration effects operate through urban *density*, not population, which depends on the ratio of two endogenous variables L_i and X_i, where X_i can be written as a function of w_i and L_i;
2. the agglomeration and congestion elasticities are spatially heterogeneous;
3. there is an outside sector (agriculture).

Of these, the third point is the most difficult to handle. I proceed in steps. First,

1. First, I show that the allocation of labor across cities L_i, holding L_a fixed, exists and is unique given the theorem of Allen et al. (forthcoming).
2. Once this is established, I can define an 'urban sector welfare function' $\mathcal{W}^u$ such that,

$$L_a = \frac{v_a(L_a)^\varepsilon}{(\mathcal{W}^u(L - L_a))^\varepsilon + v_a(L_a)^\varepsilon} L$$

   and show that this equation has a unique fixed point L_a, determining the equilibrium.

Since scaling constants are irrelevant to the application of Theorem 1 of Allen et al. (forthcoming), I subsume them in constants $\tilde{A}_i$,

First, the population distribution is pinned down by,

$$L_i = \left(\frac{\tilde{A}_i w_i^{1-\kappa_i} L_i^{-\kappa_i} P_i^{-\alpha}}{\mathcal{W}} \right)^\varepsilon$$

where $\mathcal{W}$ is the equilibrium expected welfare of agents. The price index remains,

$$P_i = \left(\sum_j \tau_{ji} (w_j / Z_j)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

And finally goods market clearing is,

$$w_i L_i = \sum_j \tau_{ij}^{1-\sigma} (w_i / Z_i)^{1-\sigma} P_j^{\sigma-1} \hat{\alpha} w_j L_j + \tau_{ia}^{1-\sigma} (w_i / Z_i)^{1-\sigma} P_a^{\sigma-1} \alpha w_a L_a$$

Finally,

$$Z_i = \bar{Z}_i (L_i / \pi X_i^2)^\zeta$$

Since,

$$\pi X_i^2 = \tilde{Z}_i^X (w_i L_i)^{\frac{\rho_i}{1+\rho_i}}$$

I define $\omega_i = \frac{\rho_i}{1+\rho_i}$ so that in equilibrium,

$$Z_i = \bar{Z}_i L_i^{\zeta} (w_i L_i)^{-\omega_i \zeta} = \bar{Z}_i L_i^{\zeta(1-\omega_i)} w_i^{-\omega_i \zeta}$$

so when land supply is perfectly inelastic $\omega_i = 0$ and the returns to population are the same as the returns to density.

Using the population equation, we can write,

$$P_i = \left(\frac{\tilde{A}_i w_i^{1-\kappa_i} L_i^{-\kappa_i}}{\mathcal{W}} L_i^{-1/\varepsilon} \right)^{1/\alpha} = \hat{A}_i (\mathcal{W})^{-1/\alpha} w_i^{\frac{1-\kappa_i}{\alpha}} L_i^{-\frac{1}{\alpha}(\kappa_i + 1/\epsilon)}$$

Now write market clearing as,

$$\begin{aligned}
w_i L_i &= \sum_j \tau_{ij}^{1-\sigma} \left(w_i / Z_i\right)^{1-\sigma} P_j^{\sigma-1} \hat{\alpha} w_j L_j + \tau_{ia}^{1-\sigma} \left(w_i / Z_i\right)^{1-\sigma} P_a^{\sigma-1} \alpha w_a L_a \\
&= \sum_j \tau_{ij}^{1-\sigma} \left(w_i / Z_i\right)^{1-\sigma} \left(\hat{A}_i (\mathcal{W})^{-1/\alpha} w_j^{\frac{1-\kappa_i}{\alpha}} L_j^{-\frac{1}{\alpha}(\kappa_i - 1/\epsilon)} \right)^{\sigma-1} \alpha w_j L_j \\
&+ \tau_{ia}^{1-\sigma} \left(w_i / Z_i\right)^{1-\sigma} P_a^{\sigma-1} \alpha w_a L_a \\
&= \sum_j \tau_{ij}^{1-\sigma} \left(w_i / \bar{Z}_i L_i^{\zeta(1-\omega_i)} w_i^{-\omega_i \zeta} \right)^{1-\sigma} \left(\hat{A}_i (\mathcal{W})^{-1/\alpha} w_j^{\frac{1-\kappa_i}{\alpha}} L_j^{-\frac{1}{\alpha}(\kappa_i + 1/\epsilon)} \right)^{\sigma-1} \alpha w_j L_j \\
&+ \tau_{ia}^{1-\sigma} \left(w_i / Z_i\right)^{1-\sigma} P_a^{\sigma-1} \alpha w_a L_a \\
w_i^{\sigma + (\sigma-1)\omega_i \zeta} L_i^{1+(1-\sigma)\zeta(1-\omega_i)} &= \mathcal{W}^{\frac{1-\sigma}{\alpha}} \sum_j \tau_{ij}^{1-\sigma} \left(\bar{Z}_i\right)^{\sigma-1} \hat{A}_i^{\sigma-1} \alpha w_j^{1+(\sigma-1)\frac{1-\kappa_i}{\alpha}} L_j^{1+(1-\sigma)\frac{\kappa_i + 1/\epsilon}{\alpha}} \\
&+ \tau_{ia}^{1-\sigma} \left(Z_i\right)^{1-\sigma} P_a^{\sigma-1} \alpha w_a L_a
\end{aligned}$$

The price index equation is,

$$\begin{aligned}
\hat{A}_i(\mathcal{W})^{-1/\alpha} w_i^{\frac{1-\kappa_i}{\alpha}} L_i^{-\frac{1}{\alpha}(\kappa_i+1/\epsilon)} &= \left(\sum_j \tau_{ji} (w_j/\tilde{Z}_j)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \\
w_i^{\frac{1-\kappa_i}{\alpha}} L_i^{-\frac{1}{\alpha}(\kappa_i+1/\epsilon)} &= \hat{A}_i^{-1}(\mathcal{W})^{1/\alpha} \left(\sum_j \tau_{ji} (w_j/Z_j)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \\
w_i^{(1-\sigma)\frac{1-\kappa_i}{\alpha}} L_i^{-(1-\sigma)\frac{1}{\alpha}(\kappa_i+1/\epsilon)} &= \hat{A}_i^{\sigma-1}(\mathcal{W})^{\frac{1-\sigma}{\alpha}} \left(\sum_j \tau_{ji} (w_j/Z_j)^{1-\sigma} \right) \\
w_i^{(1-\sigma)\frac{1-\kappa_i}{\alpha}} L_i^{-(1-\sigma)\frac{1}{\alpha}(\kappa_i+1/\epsilon)} &= \hat{A}_i^{\sigma-1}(\mathcal{W})^{\frac{1-\sigma}{\alpha}} \left(\sum_j \tau_{ji} (w_j/\bar{Z}_j L_j^{\zeta(1-\omega_j)} w_j^{-\omega_j\zeta})^{1-\sigma} \right) \\
w_i^{(1-\sigma)\frac{1-\kappa_i}{\alpha}} L_i^{-(1-\sigma)\frac{1}{\alpha}(\kappa_i+1/\epsilon)} &= \hat{A}_i^{\sigma-1}(\mathcal{W})^{\frac{1-\sigma}{\alpha}} \left(\sum_j \tau_{ji} \bar{Z}_j^{\sigma-1} L_j^{(\sigma-1)\zeta(1-\omega_j)} w_j^{(1-\sigma)(1-\omega_i\zeta)} \right)
\end{aligned}$$

To follow the proof technique of Allen et al. (forthcoming), I must define matrices whose entries are the righthandside (B) and lefthandside (Γ) coefficients. However, the matrix of lefthandside coefficients depends on the size agricultural sector in each market clearing equation. Holding the size of the agricultural sector fixed, these elasticities are bounded above by the elasticities of the righthandside when not accounting for agriculture. E.g., recall that if,

$$y = x^a + b \implies \frac{d\log y}{d\log x} = a\frac{x^a}{x^a+b}.$$

As it is required to apply the sup norm elementwise to the resulting matrices $B\Gamma^{-1}$, I can replace the i-specific elasticities with their upper bounds.

The matrix of righthandside coefficients,

$$B = \begin{pmatrix} 1+(\sigma-1)\frac{1-\kappa_i}{\alpha} & 1+(1-\sigma)\frac{\kappa_i+1/\epsilon}{\alpha} \\ (1-\sigma)(1-\omega_i\zeta) & -(1-\sigma)\zeta(1-\omega_i) \end{pmatrix}$$

And for the lefthandside coefficients,

$$\Gamma = \begin{pmatrix} \sigma+(\sigma-1)\omega_i\zeta & 1+(1-\sigma)\zeta(1-\omega_i) \\ (1-\sigma)\frac{1-\kappa_i}{\alpha} & -(1-\sigma)\frac{\kappa_i+1/\epsilon}{\alpha} \end{pmatrix}$$

To apply the theorem of Allen et al. (forthcoming), we must check whether the spectral radius of $\sup_i |B\Gamma^{-1}|$ is bounded above by 1. Analytically, solving for the eigenvalues does not yield a simple parameter restriction. Instead, the condition can be numerically checked by taking the $\sup_i$ of the elasticities and then checking the condition numerically. Figure A9 displays the results of this check. It would take very large values of the agglomeration parameter ζ to give rise to an equilibrium that is not unique, and the quantified model is far from the cutoff.

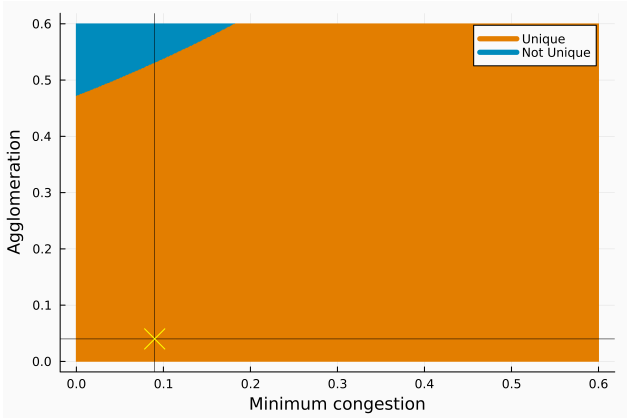


Figure A9: Check that the calibrated model has a unique equilibrium, holding the size of L_a fixed. 'Min Congestion' refers to $\min_i \kappa_i$ while agglomeration refers to ζ.

Now that it is established that the equilibrium is unique holding fixed L_a, then the value of expected utility in the urban sector must have a unique value. I must show that there is a unique L_a that solves,

$$L_a = \frac{v_a(L_a)^\varepsilon}{(\mathcal{W}^u(L - L_a))^\varepsilon + v_a(L_a)^\varepsilon} L \tag{22}$$

where,

$$\mathcal{W}^u = \left(\sum_{i=1}^{I} (v_i(L_i))^\varepsilon \right)^{1/\varepsilon}.$$

There is a unique solution to (22) provided the righthandside slopes down in L_a. In elasticity form, the slope of the righthandside is,

$$\frac{d \log \text{RHS of (22)}}{d \log L_a} = \left(1 - \frac{L_a}{L}\right) \varepsilon \left(\frac{d \log v_a}{d \log L_a} - \frac{d \log \mathcal{W}^u}{d \log L_a} \right).$$

For the term $\frac{d \log \mathcal{W}^u}{d \log L_a}$ is must be bounded above by $\max_i \kappa_i$ while, $\frac{d \log v_a}{d \log L_a} = \beta + (1-\beta)\mu = 0.25 + 0.75 \times 0.7 = 0.775$. Numerically, the 99.9th percentile across κ_i is 0.4, so the equation is downward sloping in L_a.

C Data descriptions

Table A8 lists all data sources used. All data processing was performed in Google Earth Engine before exporting to CSV for analysis. Maps were produced in R.

Table A8: Data sources used in this project.

Data	Source	Notes
Cities		
Built volume within cities	Global Human Settlement Layer (GHSL)	10m×10m pixel extent aggregated to 100m×100m using Google Earth Engine
Global extent of cities	Urban centers database (GHSL) (Florczyk et al., 2019)	City polygons, names, countries, area, and GDP (Kummu et al., 2018) and population (GPW)
Cities' central business districts	Google Maps	Unqueriable cities assigned polygon centroids as the CBD
Cities' elevation and slope	NASADEM	30m resolution
Land cover	Space Agency WorldCover (Zanaga et al., 2022)	10m resolution
City soil content	OpenLandMap	250m resolution – 1m band used for all products
High-resolution nightlights	VIIRS Nighttime Day/Night Band Composites	463.83m resolution, 2015 annual average DNB radiance values used.
City population	Gridded Population of the World (GPW)	
City GDP (not nightlights-extrapolated)	Kummu et al. (2018)	Used for GDP extrapolation outside U.S./EU countries
Lower-resolution nightlights	DMSP-OLS	927.67m resolution, Stable lights used. 1992 used for 1990, 2013 used for 2015.
Regulatory burden	Share of warehouse construction costs allocated to permits, World Bank 'Doing Business' survey	

Continued on next page

Table A8 – continued from previous page

Data	Source	Notes
Cities' transportation infrastructure	Average values of roadway characteristics (weighted by road length) from OSMnx (Boeing, 2024)	Variables like share of roads paved, average road width, etc. Values taken within 10 km of the CBD or the maximum radius of the city, whichever is smaller.
Average travel distance between cities	OpenStreetMaps	Computed on the road network in km
Countries		
Agricultural share value-added of GDP	World Development Indicators (WDI)	I treat Taiwan as a separate entity than China, and rely on data from Our World In Data for its aggregate statistics, and assign it the agricultural share of value-added from South Korea
Share of population employed in agriculture	WDI	
National GDP, USD	WDI	
PPP conversion factor	WDI	
Historical urbanized population share	WDI	
Cropland extent	Thenkabail et al. (2021)	
Cropland temperature and global temperature	Berkeley Earth Surface Temperatures (BEST)	Annual averages over cropland (computed) and global temperature series
Average floorspace prices	Liotta et al. (2022)	See Appendix C.3
U.S. cities – for validation exercises		
U.S. cities floorspace supply elasticities	Baum-Snow and Han (2024)	
U.S. cities land supply elasticities	Saiz (2010)	
Floorspace price gradients	Zillow ZHVI	2-bedroom home data
Employment, commuting, and residency patterns	Longitudinal Employer Household Dynamics Longitudinal Origin Destination Employment Survey (LEHD-LODES)	

I now detail the data definitions used in the paper.

C.1 Developable land

Following Saiz (2010) and Harari (2020), I define land as suitable for development if its slope does not exceed 15%, as measured by the NASADEM elevation data, which is available at the 30m resolution. I combining this with the European Space Agency WorldCover data classifications. I refine the definition of land suitable for development (developable land) if the land cover classification does not classify the land as bodies of water, permanent ice, or permanent wetlands. For each 100m$\times$100m pixel within a city, I compute the share of land that is developable using this definition.

C.2 City GDP

The city GDP data provided by Kummu et al. (2018) matched to the cities in the data is measured with substantial error. This is because the data product takes subnational measures of output, temporarily extrapolates them, and then population-weights them using the GPW to distribute that GDP over space. Consequently, two cities in the same administrative unit in which GDP data is recorded do not differ in terms of GDP per capita when appropriately using the GPW. While this is not a problem in the developed world, where subnational output data often corresponds to cities' political boundaries and is recorded annually, it is a substantial issue in the developing world where such data are available at irregular frequencies in time and at very coarse spatial resolutions. For example, in their 2015 data, for many countries in sub-Saharan Africa, GDP data is only available at the country level, meaning that income variation across cities within a nation in these data are driven solely by the population data. Using this data in regression-based estimation would bias all estimates downwards in developing nations!

To avoid spatially heterogeneous attenuation bias and provide a common definition of city income across all cities, I instead rely on a high-resolution nighttime luminosity series from the VIIRS satellite data. This data improves over the DMSP-OLS nightlights data in that it is available at a more granular spatial resolution and does not suffer the well-known issues due to top and bottom coding in that time series. Moreover, it is available for post-2012 years, while the DMSP-OLS is not. It is recommended as a superior product for inferring local economic activity in the development literature. Indeed, Gibson et al. (2021) finds that the"city lights-GDP relationship is twice as noisy with DMSP data than with VIIRS data." To map average radiance values over an urban boundary to city income, I run a regression of the form,

$$\log GDP_i = \eta \log \text{Average city radiance}_i + \xi_n + e_i$$

where ξ_i is a country fixed effect. I restrict the sample to only cities in the North America and Europe. I then form predicted values out of sample,

$$\widehat{GDP}_i = \exp(\hat{\eta} \log \text{Average city radiance}_i)$$

which I then rescale at the country level so that $\sum_i \widehat{GDP}_i$ is equal to the share of national income that is not agricultural income. As I include country or city fixed effects in all regressions that rely on city GDP, this scaling does not matter for estimation. Instead, its purpose is to match the market clearing constraints implied by the model when quantifying the model. This procedure is similar to the procedure used in Ducruet et al. (2024) when extrapolating port cities' GDP. Table A9 reports the results of this regression.

	(1)
Log Average Radiance	0.922
	(0.015)
R-squared	0.833
R-squared (within)	0.785
Observations	1,120

Table A9: The relationship between log VIIRS average radiance and city GDP from Kummu et al. (2018) for cities in North America and Europe.

C.3 Average floorspace prices

When quantifying the model, $\hat{A}_i$ is identified alone off the data on population, income, the intercity trade network, and the urban cost elasticity parameters, alongside other model elasticities. The terms Z_i^H/ψ_n^H and Z_i^X/ψ_n^X are identified using the urban cost elasticities and city income, floorspace, and area.

However, to compute the level of Z_i^H or Z_i^X, to make international comparisons in building productivity, I require floorspace prices to disentangle Z_i^H and Z_i^H from their quality adjustment terms, ψ_n^H and ψ_n^X, which are assumed constant within a nation. This does not matter for counterfactuals, since I hold ψ^H and ψ^X terms constant and they enter the indirect utility function multiplicatively.

Built volume and income data alone cannot separate physical productivity in floorspace Z_i^H from its quality. For example, the product Z_i^H/ψ^H appears as large as it is in New York as it is in Lagos, Nigeria. However, the quality of floorspace varies substantially across these cities; in short the term ψ^H is doing a lot of work equating 'perceived TFP' from actual physical productivity. To net out these quality adjustment terms requires floorspace prices which are not available at the city or even national level.

However, Liotta et al. (2022) recover data on floorspace prices for around 192 cities in 49 countries. I use these data to construct average floorspace prices, for all countries around the world, which I predict using machine learning, and then use the model's market clearing constraints to recover ψ^H.

I take the average price of floorspace within a country using their regression-adjusted (e.g., accounting for number of bedrooms) estimates at the country-currency-year level. I convert local currencies to USD using the World Bank's PPP deflator. I then use the U.S. CPI to adjust prices in different years to account for U.S. inflation, to put all values in 2015 dollars, and then average these again, weighing by the number of observations in a country-year-currency cell, to a country average.

This procedure gives me 49 observations of average floorspace prices around the world. For five countries, the units provided in the data appear off, they are several orders of magnitude higher or lower than othe reported observations, even when accounting for a nation's GDP per capita. These are Cote d'Ivoire, Morocco, Slovenia, Argentina, and Croatia. This still leaves me with ample coverage of developing and middle-income nations: the data include Ethiopia, India, and Indonesia, for example.

I then train a random forest learner to predict average floorspace prices using the World Bank's PPP deflator, log GDP/capita, the size of a country, and the Z-score of the 'building quality' index from the World Bank's 'Doing Business' survey. 67% of the variation in observed price is captured by the predicted price. Figure

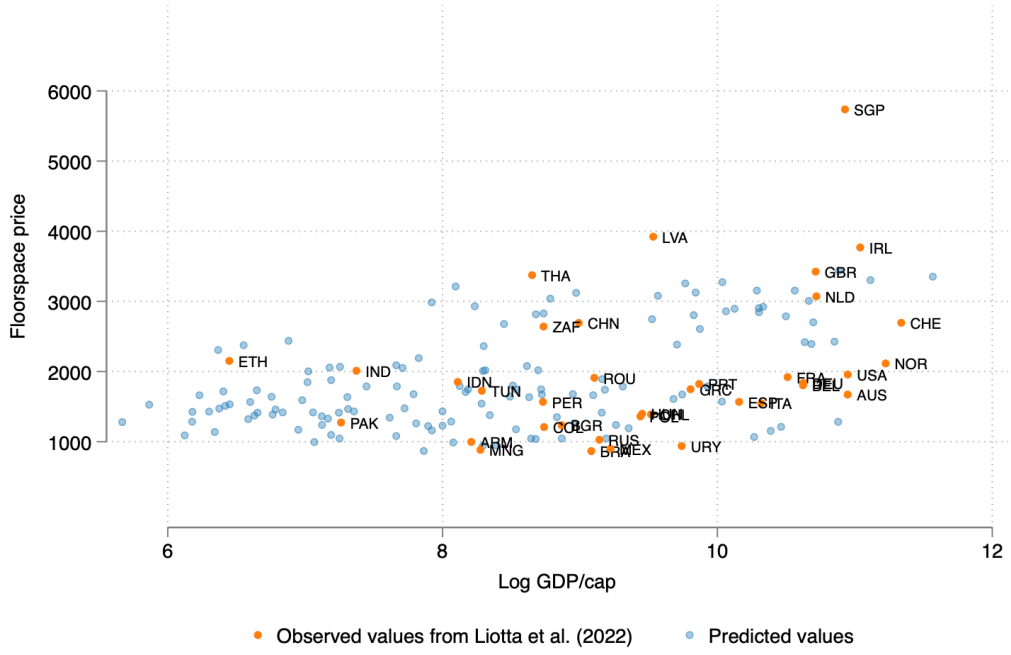


Figure A10: Predicted and observed floorspace prices using data from Liotta et al. (2022) and a random forest learner.

A10. The values look reasonable. For the U.S., the average price per square meter is $1,960, or about $182 per square foot.

D Estimation

This section details the estimation of the parameters of the urban technology: the commuting cost elasticity (τ_i), the floorspace supply elasticity (γ_i), and the land supply elasticity (ρ_i).

D.1 Estimating cities' skyline slopes, $\tau_i\gamma_i/\beta$

First, for each city i, I estimate,

$$H_i(x,\phi) = \exp\left(\frac{\tau_i\gamma_i}{\beta}\log x + \xi_{i,\phi} + t_i(x)\right)$$

with a Poisson psuedo-maximum likelihood (PPML) estimator. The error term $t_i(x)$ represents measurement error in $H_i(x,\phi)$, and I assume $\log \mathbb{E}[t_i(x) \mid \log x, \xi_{i,\phi}] = 0$, where $\xi_{i,\phi}$ are city-by-polar angle fixed effects. I include 8 polar angle fixed effects for each city. Within each city, I bin observations into 100 bins of x and reweight observations so that the weights are inversely proportional to the number of observations per bin, so that there are the same effective number of observations at each distance horizon, undoing the mechanical rise in observations as x increases due to the simple geometry of circles.

Calling $\hat{\theta}^{PPML} = -\widehat{\frac{\tau_i\gamma_i}{\beta}}$, some estimates of $\hat{\theta} < 0$ corresponding to an upward sloping skyline gradient, inconsistent with the model. However, positive and near-zero estimates are on average estimated with less precision than the majority of estimates around the mean, see the left panel of Figure A11. Consequently, to ensure that $\hat{\theta} > 0$ for all estimates, I employ an empirical Bayes' shrinkage estimator,

$$\begin{aligned} \hat{\theta}_i^{PPML} \mid \theta_i &\sim N(\theta_i, \sigma_i) \\ \theta_i &\sim N_{(-\infty,0)}(\theta_n, \sigma_n) \end{aligned} \tag{23}$$

Where $N_{(-\infty,0)}(\theta_n, \sigma_n)$ is a truncated normal distribution over $(-\infty, 0)$ with mean θ_n for nation n, with a country-specific standard error, as well.

The empirical Bayes' estimates are $\hat{\theta}_i^{EB} = \mathbb{E}[\theta_i \mid \hat{\theta}_i^{PPML}]$, given the model (23). As a truncated normal prior is conjugate with a normal likelihood, estimation of $\hat{\theta}_i^{EB}$ is straightforward, following Morris (1983). Using properties of the normal distribution,

$$\begin{aligned} \theta_i &= \frac{\hat{\sigma}_i^{-2}}{\hat{\sigma_i}^{-2} + \hat{\sigma_n}^{-2}}\hat{\theta}^{PPML} + \frac{\hat{\sigma}_n^{-2}}{\hat{\sigma}_i^{-2} + \hat{\sigma}_n^{-2}}\hat{\theta}_n \\ B_i^2 &= \hat{\sigma}_i^2\frac{\hat{\sigma}_n^{-2}}{\hat{\sigma}_i^{-2} + \hat{\sigma}_n^{-2}} \end{aligned}$$

where B_i^2 is the variance of the posterior distribution, $f(\theta_i \mid \hat{\theta}_i^{PPML})$. The mean of the posterior distribution is given by the Inverse Mills ratio. This object provides an analytic expression for the mean of the truncated normal distribution given its parameters,

$$\hat{\theta}_i^{EB} = \mathbb{E}[\theta_i \mid \hat{\theta}_i^{PPML}] = \theta_i - B_i\frac{\phi(-\theta_i/B_i)}{\Phi(-\theta/B_i)} \tag{24}$$

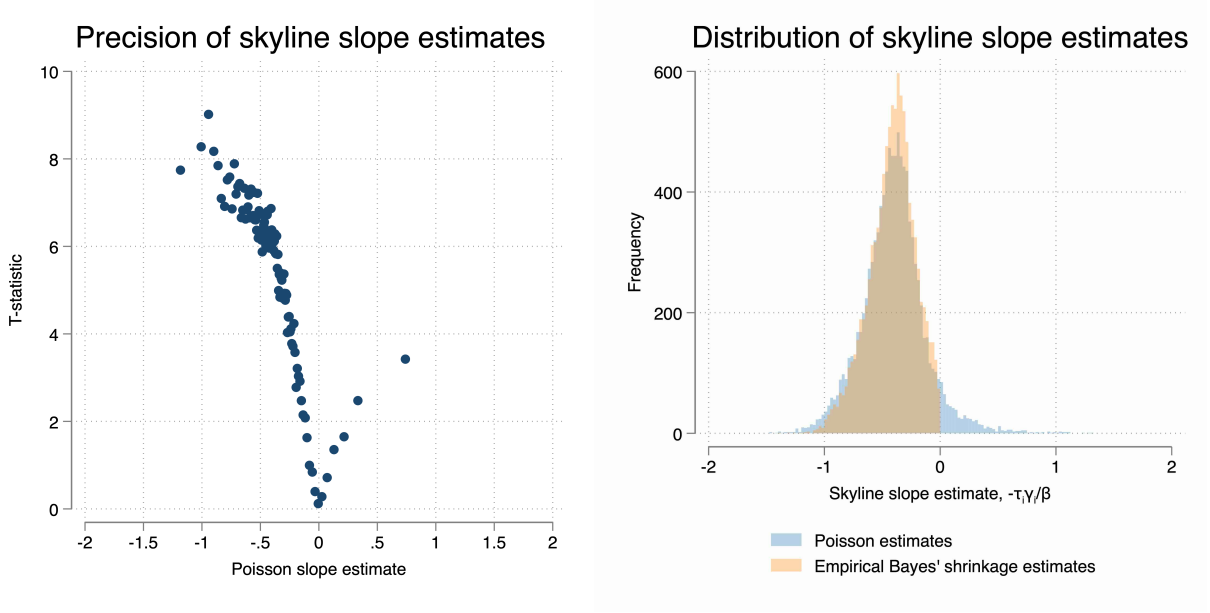


Figure A11: Left: precision of $\hat{\theta}$ (measured by a T-statistic), versus estimated values of θ.

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the pdf and cdf of the standard normal distribution. This ensures posterior estimates are contained on $(-\infty, 0)$.

Following Morris (1983), the variance σ_n^2 is estimated by solving the fixed point of,

$$\hat{\sigma}_n^2 = \frac{\sum_{i \in n} W_i \left((\theta_i^{PPML} - \bar{\theta}_n^{PPML})^2 - \sigma_i^2 \right)}{\sum_{i \in n} W_i}$$

$$\bar{\theta}_n^{PPML} = \frac{\sum_{i \in n} W_i \hat{\theta}_i^{PPML}}{\sum_{i \in n} W_i}$$

$$W_i = \frac{1}{\hat{\sigma}_i^2 + \hat{\sigma}_n^2}.$$

To solve the fixed point, I initialize $\hat{\sigma}_n$ at 1, then,

Step 1. form W_i

Step 2. construct $\bar{\theta}_n^{PPML}$

Step 3. estimate $\hat{\sigma}_n^2$

and repeat steps 1-3 until $\hat{\sigma}_n^2$ converges. I do this for each nation n in the data.

Once $\hat{\sigma}_n^2$, I construct B_i and form the estimated mean of the posterior conditional on $\hat{\theta}_i^{PPML}$ using equation (24).

The right panel of A11 shows the density of the Poisson estimates (blue) and the distribution of the empirical Bayes' estimates (orange). The empirical Bayes' estimates mainly shrink observations of $\theta_i > 0$ towards the mean, such that they fall below zero.

D.2 Estimating the floorspace supply elasticity, γ_i

Implied by the model, I begin with the estimating equation,

$$\log \frac{\bar{H}_i}{\pi X_i^2} = \frac{\gamma_i}{1+\gamma_i}\left(\log w_i + \log \frac{L_i}{\pi X_i^2}\right) + a_i, \tag{25}$$

where the error term a_i is equal to,

$$a_i = \frac{1}{1+\gamma_i}\log\left(\frac{Z_i^H}{\psi^H}\right) + \log\left(\frac{\left(1-\frac{\tau_i}{2}\frac{1+\gamma_i}{\beta}\right)^{\frac{\gamma_i}{1+\gamma_i}}}{1-\frac{\tau_i}{2}\frac{\gamma_i}{\beta}}\right).$$

The error term reflects the endogeneity issue, that Z_i^H is observed but is correlated with city size L_i in equilibrium. The regressor w_i is additionally endogenous, as,

$$w_i = p_i \underbrace{\bar{Z}_i^y \left(\frac{L_i}{\pi X_i^2}\right)^{\zeta}}_{\equiv Z_i^y},$$

i.e., it is directly influenced by a city's population and area through agglomeration effects, and indirectly, through prices. Price endogeneity stems from the fact that goods demand slopes down. Larger cities produce more, moving down the demand curve and generating lower equilibrium prices. However, through the structure of the model, this price effects can be controlled for by recovering the exogenous component of wages, $\bar{Z}_i^y$.

To handle the endogeneity of w_i, I instrument for it using $\bar{Z}_i^y$.

D.2.1 Constructing model-generated productivity instruments, $\bar{Z}_i^y$

Using goods market clearing, these productivity terms are effectively observed without any knowledge of γ_i. By goods market clearing,

$$w_i L_i = \sum_{j=a,1,\ldots,N_n} \left(\frac{p_{ij}}{P_j}\right)^{1-\sigma} \alpha w_j L_j.$$

Plugging in for prices, where $p_{ij} = \delta_{ij} p_i$,

$$w_i L_i = \sum_{j=a,1,\ldots,N_n} \delta_{ij}^{1-\sigma} \left(\frac{w_i/Z_i^y}{\left(\sum_k \left(\delta_{kj}(w_k/Z_k^y)\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}}\right)^{1-\sigma} \alpha w_j L_j. \tag{26}$$

Using w_i and L_i in the data, alongside $\{\delta_{ij}\}$, allows me to iteratively solve for Z_i^y by solving the fixed point of the system (26), and then remove the endogenous component under the assumption that $\zeta = 0.04$.

While these instruments allow me to recover the effect of w_i on average height, an issue still remains: I only

have the cross-section on building height, which means I cannot estimate city-specific γ_i. To handle this issue, I model γ_i as a function of observables.

D.2.2 Projecting γ_i onto observables

I project γ_i onto city specific observables by assuming,

$$\frac{\gamma_i}{1+\gamma_i} = G_i'\Gamma$$

where G_i is a vector of city- or nation-level observables and Γ is a vector of coefficient loadings. Consequently, I estimate,

$$\hat{\gamma}_i = \frac{G_i'\hat{\Gamma}}{1 - G_i'\hat{\Gamma}}.$$

I estimate Γ by estimating,

$$\log \frac{\bar{H}_i}{\pi X_i^2} = \sum_k \left(\Gamma^k G_i^k \times \log w_i + \Delta^k G_i^k \log \frac{L_i}{\pi X_i^2} + \Theta^k G_i^k \right) + \xi_n + \tilde{a}_i, \tag{27}$$

where ξ_n is a country fixed effect, and $\tilde{a}_i$ is the error term. G_i^k is the value of kth observable for city i, and I assume the $k = 1$th entry is a constant. I estimate (27) by two-stage least squares, using $\log \bar{Z}_i^y$ and its interactions with G_i^k as instruments.

In addition to a constant, I include in G_i the following variables: the share of developable land within 10km of the CBD, the average slope of a city, the average elevation of a city, the average clay content in the soil at 1m, the average water content in the soil at 1m, the average soil density at 1m, the average sand content in the soil at 1m, and the average share of construction costs allocated to permits from the World Bank 'Doing Business' survey.[45]

For each variable, I use an arctangent transformation. In particular, I standardize the variable by substracting the mean and dividing by the standard deviation across cities to construct $\tilde{G}_i^k$ so that,

$$G_i^k = \arctan\left(\frac{\pi}{2}\tilde{G}_i^k\right).$$

This ensures that values of $G_i^k \approx \tilde{G}_i^k$ within one standard deviation, but outliers have a log-like transformation applied. Multiplying by $\pi/2$ ensures that transformation maps 1 to 1 and -1 to -1, and the function is approximately linear in that interval. The reason I do this is that this is purely a predictive exercise, so I am not interested in recovering the causal effect of a variable on the floorspace supply elasticity. Instead, I am more concerned with the predicted value of γ_i for outliers in the data. Moreover, because constructing γ_i requires a nonlinear transformation, I must ensure that values $\sum_k G_i'\hat{\Gamma}^k \in (0,1)$ as implied by theory, which allows $\hat{\gamma}_i$ to be well-defined.

The left column of Table A10 reports results from this regression.

[45] See: https://archive.doingbusiness.org/en/data/exploretopics/dealing-with-construction-permits

D.3 Estimating the land supply elasticity, ρ_i

The model implies the relationship,

$$\log \pi X_i^2 = \frac{\rho_i}{1+\rho_i} \log w_i L_i + b_i$$

where,

$$b_i = \log \pi + \log \frac{Z_i^X}{\psi^X}$$

reflecting the endogeneity of the regressor: higher Z_i^X cities are larger in equilibrium.

Unlike city average height, which is only available in the cross-section, I have the time series on area for years t =1990, 2000, and 2015 (1975 is available for area, but not income).

To identify ρ_i, I rely on the time series. Identification comes from the idea that, that the intercept of the land supply curve Z_i^X is driven by time-invariant factors, and common shifts do to changes in commonly-available building technology. My estimating equation is,

$$\log \text{Area}_{it} = \sum_k \left(\Omega^k G_i^k \times \log GDP_{it} \right) + \xi_i + \xi_{rt} + \tilde{b}_i, \tag{28}$$

where ξ_i, and ξ_{rt} are city and region-year fixed effects and Area is the sample analog of πX_i^2 and GDP_{it} is the sample analog of $w_{it} L_{it}$. In essence, the idea is these fixed effects hold the land supply curve fixed so that residual movement in GDP_{it} and Area_{it} traces out movement along the land supply curve from shifts in demand. Ω^k is a vector of coefficients so that,

$$\hat{\rho}_i = \frac{G_i' \hat{\Omega}}{1 - G_i' \hat{\Omega}}.$$

However an issue remains: how to construct GDP_{it}. In the estimation sample used for estimating γ_i and used in the quantification of the model, city GDP is constructed using a high-resolution nightlights series; see Section C.2. This data (the VIIRS) are unavailable historically (the VIIRS satellite was launched in 2011).

This leaves me with two potential sources of historical city GDP data: that from Kummu et al. (2018) and extrapolated data from the DMSP-OLS nightlights series. The data from Kummu et al. (2018) are downscaled subnational accounts data that are population weighted using gridded population of the world. Cities that fall in the same spatial unit in which GDP data is available only differ in income according to their population. In developing nations, the data is very spatially coarse and therefore city income is measured with substantial error. In the DMSP-OLS series, the resolution of the data is quite coarse, there is substantial top- and bottom-coding, and the pixel values are not well-calibrated so as to make year-over-year comparisons. In short, both of these potential income time series are measured with considerable error. Measurement error in the regressor of interest would bias the coefficient downwards.

To overcome this, I treat the measurement error in the Kummu et al. (2018) data as classical and instrument for the reported values using the DMSP-OLS nightlights. To form instruments, I use the log of the sum of pixel values in the DMSP-OLS series within each UCDB city shapefile for years 1992 (associated with 1990), 2000, and 2012 (associated with 2015). I interact the instrument with covariates in G_i to form a full set of instruments.

Column (2) of Appendix Table A10 reports the results of this regression. The share of developable land available is the largest driver of differences in the land supply elasticity across space.

	(1) Log average height	(2) Log area
Constant	0.447	0.371
	(0.0307)	(0.0722)
Developable share	-0.0198	0.262
	(0.0322)	(0.0440)
Slope	-0.0109	0.0525
	(0.00599)	(0.00787)
Elevation	-0.0300	0.0393
	(0.00621)	(0.00815)
Clay	-0.0484	-0.0195
	(0.00685)	(0.00961)
Soil water	0.0370	-0.0303
	(0.00817)	(0.0109)
Soil density	0.0223	-0.0122
	(0.00681)	(0.00981)
Sand	-0.0574	-0.0376
	(0.00711)	(0.00951)
Permit share of warehouse cost	-0.0166	0.0476
	(0.0131)	(0.0215)
Country fixed effects	✓	✓
City fixed effects		✓
N	9,132	26,574
Kleibergen-Paap F-statistic	201.0	22.59

Table A10: Estimation results. Column (1) reports estimates of Γ^k from estimating (27). Column (2) reports estimates of Ω^k from estimating (28). Standard errors clustered at the city level in parentheses.

E Model validation

The model makes several assumptions that I seek to validate in a context where data is abundant: the United States. First, I test the model's strongest assumption: urban monocentricity. To test for monocentricity, I examine how several model predictions bare out in data on residency, employment, and commuting patterns and housing prices in U.S. data. I then examine how well my measures of τ_i, γ_i and ρ_i predict prices and price and population gradients within U.S. cities. Finally, I examine whether my estimates of these elasticities are consistent with other estimates in the literature.

First, I test whether commute distance increases in a resident's distance to the CBD, whether the residency gradient is decreasing in distance to the CBD, whether employment is primarily concentrated in the CBD, and whether floorspace prices fall in distance to the CBD. On all four counts, the data hold up to the assumptions of the monocentric city model. I rely on the Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics data (LODES) for commuting patterns and the spatial distribution of residency and employment in U.S. cities, and Zillow's Home Value Index (ZHVI) for data on home prices at the zipcode (zip5) level. The LODES data provide census block-to-block tabulations of commuting flows, which I aggregate to the census tract level in 2015. In the ZHVI, I examine the average price of a 2-bedroom home in 2015.

	Average commute length		Density gradient		
	(1) Residents	(2) Workers	(3) Residents	(4) Workers	(5) Resident gradient
Distance to CBD (km)	0.203	0.315			
	(0.033)	(0.069)			
Log distance to CBD			-0.765	-1.258	
			(0.022)	(0.023)	
τ_i					0.680
					(0.536)
γ_i					0.103
					(0.081)
Log population					-0.132
					(0.014)
R-squared	0.837	0.541	0.358	0.371	0.301
Within R-squared	0.017	0.015	0.208	0.272	
City FE	✓	✓	✓	✓	
WLS					✓
Observations	37,990,129	52,253,072	15,455	15,542	322

Table A11: Regressions of average commute length (as the crow flies) and the population and employment density gradient as a function of distance to a city's CBD (columns 1-4), controlling for city fixed effects. Robust standard errors in parentheses. Column 5 regresses the estimated resident density gradient against estimated parameters τ_i, γ_i, controlling for a city's population. WLS refers to weighted least squares; I reweight observations to undo the mechanical rise in observations as you leave the CBD.

	(1) Log ZHVI	(2) Price gradient	(3) Price gradient	(4) Price gradient	(5) Log ZHVI	(6) Log ZHVI
Log distance to CBD	-0.102					
	(0.036)					
τ_i		1.119	2.590	2.114	-0.968	-1.780
		(0.653)	(0.673)	(0.820)	(0.779)	(0.777)
γ_i				0.006	-0.181	-0.201
				(0.114)	(0.113)	(0.115)
ρ_i					-0.288	-0.250
					(0.086)	(0.084)
Log population				0.132		0.112
				(0.025)		(0.034)
R-squared	0.748	0.025	0.101	0.431	0.042	0.084
R-squared (within)	0.160					
City FE	✓					
WLS		✓	✓	✓		
At least 10 zipcodes			✓	✓		
Observations	3,006	259	103	103	320	320

Table A12: Column 1 reports a regression of the floorspace price gradient in U.S. cities. Columns 2-4 regress the price gradient on τ_i and γ_i in different samples. Columns 3 and 4 drop cities with too few observations ($<$10) which makes price gradient estimation very noisy. Columns 5-6 regress average floorspace prices at the city level on model parameters. Robust standard errors in parentheses.

Appendix Table A11 displays correlations between variables computed using the LODES and the distance to a city's CBD, subset on the U.S. cities I consider in the main analysis. First, I regresses the average commute length taken by a residents and workers on the distance of their home tract to their CBD, controlling for city fixed effects. A one kilometer increase in distance from the CBD increases the average commute length by a resident 0.4km (column 1, s.e. 0.03). However, the average commute length taken by a worker does not depend on the workplace's distance to the CBD, except in the core, where the average commute length is greater, consistent with longer distance commuters working in the core. These relationships are visually apparent in the data (Appendix Figure A12). I then compute the slope of the log population and log employment density gradients with respect to log distance to the CBD. Both are strongly negative. The employment density gradient is 68% steeper than the resident population density gradient. Visually, population density is well approximated by a log linear relationship in distance to the CBD (Appendix Figure A13, left panel) but the employment density gradient is substantially steeper in the core than the periphery (right panel), consistent with most productive activity occurring in the CBD.

I now turn towards price data. Appendix Table A12 shows that the average price gradient in U.S. cities in this data is -0.1 (column 1, s.e., 0.03). Recall from the model that the floorspace gradient falls at rate $-\tau_i/\beta$. The ZHVI data contains home prices for a 'typical' two-bedroom home, so floorspace is likely not fixed across space. Per person floorspace demand increases farther from the CBD as prices fall, and so the squarefootage of the average two-bedroom home likely increases in distance to the CBD, biasing the estimate of the gradient down. I compute the slope of the ZHVI price data for each city, and regress measures of each city's price gradient on my measures of τ_i and γ_i. In cities with at least ten zipcodes with price data, the coefficient on τ_i is 2.59 (column 3, s.e. 0.67). Controlling for a city's log population attenuates the estimate slightly to 2.11.

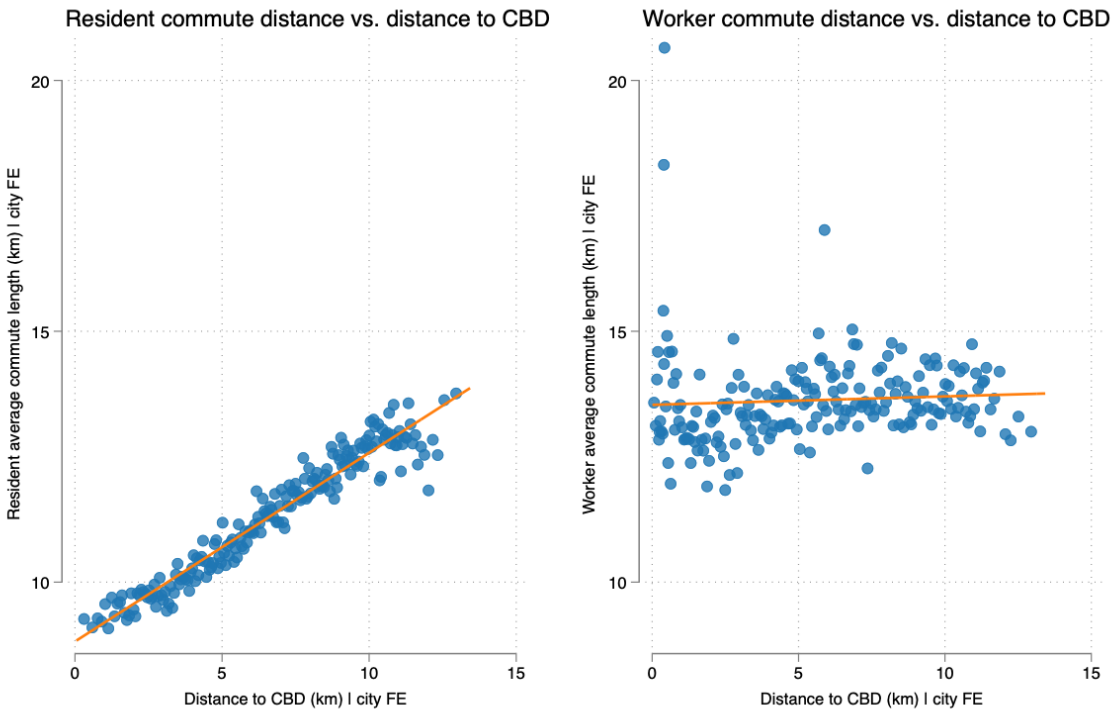


Figure A12: Conditional binscatters of resident commute distance as a function of distance to the CBD, and the average distance taken by workers at a location, as a function of distance to the CBD in the LEHD-LODES data.

Consistent with the theory, γ_i is not predictive of the price gradient (coefficient 0.01, s.e. 0.11). However in column 5 of Appendix Table A11, I regresses a city's population gradient on τ_i and γ_i and find that they both positively predict the population gradient, as the model suggests. Conditioning on a city's log population, the average log ZHVI for a city is decreasing in τ_i, γ_i, and ρ_i. This pattern is consistent with high τ_i being a negative amenity reflected in prices. Moreover, cities with more elastic floorspace supply (γ_i) and land supply (ρ_i) have lower prices in equilibrium, agreeing with the logic of the model (Appendix Table A12, column 6).

While there has been little work estimating floorspace supply elasticities globally, there is comparable work in the United States. Baum-Snow and Han (2024) computes estimates of a similar floorspace supply parameter across U.S. census tracts and aggregates them to metropolitan areas.[46] The left panel of Appendix Figure A14 plots my estimates against theirs. The correlation is weakly positive, though my estimates are twice as large. Unlike their estimates, I estimate mine in the cross-section and thus they reflect long-run floorspace supply elasticities. Their estimates reflect within-city, short-run elasticities, and capture city-specific regulations, which may explain the magnitude of my estimates and why they fail to capture some of the variation in γ_i across metro areas in the United States.[47]

[46]Their conceptual definition of a housing supply elasticity combines both my 'build up' (γ_i) and 'build out' (ρ_i) elasticities *within a census tract*, but not a metro area as a whole.

[47]Nonetheless, my estimates do capture some of the variation one might anticipate given stories of inelastic housing supply due to zoning regulations in cities in major U.S. metro areas. For example, I estimate relatively inelastic floorspace supply in Los Angeles, San Jose, and New York, compared to Chicago, which is known to have fewer housing supply restrictions.

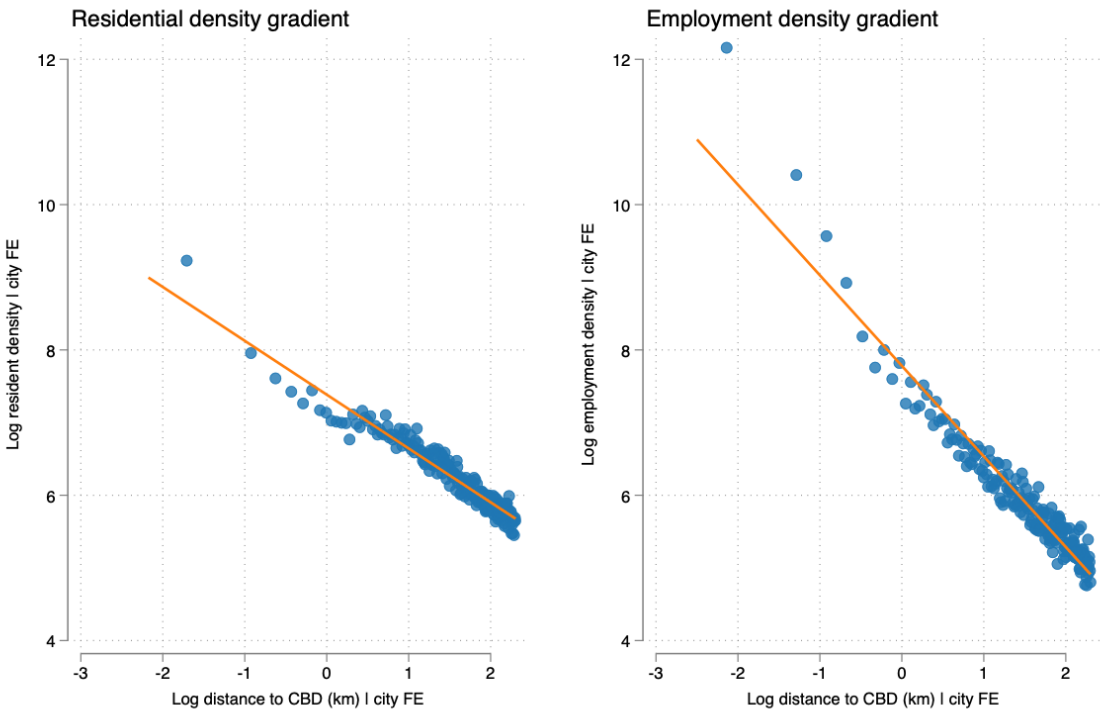


Figure A13: Conditional binscatters of population and employment density gradients within U.S. cities in the LODES data.

Similarly focused on U.S. metro areas, Saiz (2010) develops a framework in which the share of developable land directly affects the housing supply elasticity, and estimates parameters that are analogous to my land supply elasticities by regressing home values against the (instrumented) stock of houses, interacted with data on developable land and land use regulations across U.S. metro areas. My land supply estimates are strongly positively correlated with those in his study (Appendix Figure A14).

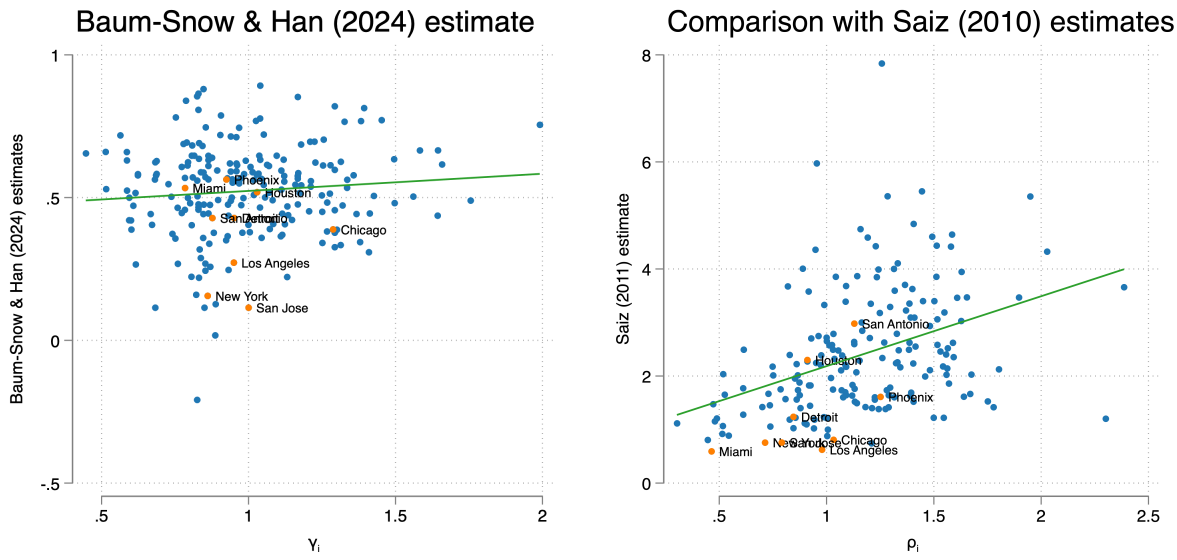


Figure A14: Left: correlation of γ_i with housing supply elasticities across U.S. metro areas from Baum-Snow and Han (2024). Right: correlation of ρ_i with land supply elasticities across U.S. metro areas from Saiz (2010)