NBER WORKING PAPER SERIES

EVALUATING TRANSPORT IMPROVEMENTS IN SPATIAL EQUILIBRIUM

Stephen J. Redding

Working Paper 33532
http://www.nber.org/papers/w33532

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2025

This paper was written while visiting Stanford University as a Trione Visiting Professor. I am grateful to Stanford University for all of its kind hospitality. This paper was commissioned for the Advances in Economics and Econometrics monograph for the World Congress of the Econometric Society in South Korea in 2025. My thinking has been influenced by the DOT-NBER Economics of Transportation in the 21st Century Research Project co-ordinated by Jim Poterba, Ed Glaeser and myself. I would like to thank Mark Bamba, Diana Furchtgott-Roth, Ed Glaeser, Benny Kleinman, Ernest Liu, Robert Hampshire, Jim Poterba, Darren Timothy, Matt Turner, and Stephan Zoepf for helpful comments and discussions. I am grateful to Bingqing Yang for excellent research assistance. Any opinions, errors and omissions are the responsibility of the author alone. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Evaluating Transport Improvements in Spatial Equilibrium
Stephen J. Redding
NBER Working Paper No. 33532
February 2025
JEL No. R30, R40, R52

ABSTRACT

The recent development of quantitative urban models provides a new set of tools for evaluating transport improvements. Conventional cost-benefit analyses are typically undertaken in partial equilibrium. In contrast, quantitative urban models characterize the spatial distribution of economic activity within cities in general equilibrium. We compare evaluations of a transport improvement using conventional cost-benefit analysis, sufficient statistics approaches based on changes in market access, and model-based counterfactuals. We show that quantitative urban models predict a reorganization of economic activity within cities in response to a transport improvement, which can lead to substantial differences between the predictions of these three approaches for large changes in transport costs.

Stephen J. Redding
Department of Economics &
School of Public and International Affairs
Princeton University
Princeton, NJ 08544
and CEPR
and also NBER
reddings@princeton.edu

1 Introduction

Transportation infrastructure investment, such as the construction of railway networks, involve some of the largest public investments undertaken. Recent years have seen the development of new theoretical models of the spatial distribution of economic activity within cities and the emergence of new sources of geographic information systems (GIS) data that record economic activity at fine spatial scales. At a time of renewed public policy interest in transportation infrastructure and rapid technological change, there remains a need for further research on the economic impact of transport infrastructure investments relative to their cost.

A major breakthrough in recent research is the development of quantitative urban models. These models are sufficiently rich to capture observed features of the data, such as many asymmetric locations and a rich geography of the transport network. Yet they remain sufficiently tractable as to permit an analytical characterization of their theoretical properties, such as the existence and uniqueness of the equilibrium. With only a small number of parameters to be estimated, these models lend themselves to transparent identification. Since they rationalize the observed distribution of economic activity in the data, they can be used to undertake counterfactuals for the impact of empirically-realistic public-policy interventions, such as the construction of railway line along a particular route.

We consider a baseline quantitative urban model following Ahlfeldt et al. (2015). The model allows for a large number of heterogeneous locations, while remaining tractable and amenable to empirical analysis. Locations can differ in productivity, amenities, the density of development (which determines the ratio of floor space to ground area), and access to the transport network. Both productivity and amenities have exogenous components that are determined by locational fundamentals (e.g., access to natural water) and endogenous components that depend on agglomeration forces. Congestion forces take the form of an inelastic supply of land and commuting costs that are increasing in travel time, as determined by the transport network.

We consider an improvement in the transport network that reduces bilateral travel times for some pairs of locations by more than for others. This reduction in commuting costs allows workers to separate their residence and workplace to take advantage of differences in productivities and amenities. Locations with high productivity relative to amenities, and good transport access to surrounding areas with high amenities, specialize as workplaces. In contrast, locations with high amenities relative to productivity, and good transport access to surrounding areas with high productivity, specialize as residences. In the presence of agglomeration forces, these exogenous differences in productivity and amenities across locations are magnified by agglomeration forces from the concentration of employment and residents across locations.

We compare evaluations of a transport improvement using conventional cost-benefit analysis, sufficient statistics approaches based on changes in market access, and model-based counterfactuals. Conventional cost-benefit analysis is typically undertaken in partial equilibrium and considers the demand for travel between locations as a function of the cost of travelling between those locations. In the special case in which the demand for travel is perfectly inelastic and the change in travel cost equals the saving in travel time multiplied by the value of time, the welfare gains from a transport improvement can be measured using conventional cost-benefit analysis as the total value of the travel time saved. More generally, this measure provides a first-order approximation to the welfare gains from a transport improvement, which can differ from its full general equilibrium impact for large changes in the transport network.

Within a class of quantitative urban models characterized by a gravity equation for commuting flows, the reorganization of economic activity in response to a transport improvement can be approximated by measures of residence and workplace market access. Residence market access captures proximity to surrounding sources of employment, while workplace market access captures proximity to surrounding sources of residents. This approximation is undertaken around an initial equilibrium with prohibitive commuting costs, and abstracts from changes in the supplies of residential and commercial floor space, and spillovers of agglomeration and dispersion forces across locations. More generally, the impact of a transport improvement depends on both market access and residuals that capture changes in the supplies of residential and commercial floor space, productivity and amenities.

We illustrate the use of quantitative urban models to evaluate a transport improvement using a numerical example of a city. By focusing on a numerical example, we consider a setting in which we know the true data generating process (DGP) and model parameters. Therefore, the data are generated according to the model, and we can examine the success of alternative approaches to approximating the true impact of the transport improvement. We show that the direct impact of a large change in the transport network on travel time can be quite misleading for the general equilibrium impact, once the reorganization of economic activity is taken into account. We show that this reorganization of economic activity can differ substantially between model-based counterfactuals and market access predictions.

This paper is related to a number of strands of existing research. First, we connect with the theoretical literature on urban economies, as synthesized in Fujita et al. (1999) and Fujita and Thisse (2002). In the canonical model of internal city structure following Alonso (1964), Muth (1969), and Mills (1967), cities are monocentric by assumption. All employment is assumed to be concentrated in a central business district (CBD). Workers face commuting costs that are increasing in the distance travelled, which gives rise to a land price gradient that is monotonically decreasing in distance from the CBD. Subsequent theoretical research relaxed the assumption that all employment is concentrated in the CBD. In important contributions, Fujita and Ogawa (1982) consider the case of a one-dimensional city along the real line, while Lucas and Rossi-Hansberg (2002) analyze a perfectly symmetric circular city. Within these frameworks, the location of both employment and residents are endogenously determined, and non-monocentric patterns of economic activity can emerge in equilibrium.

Second, we contribute to recent research on quantitative spatial economics. Within this line of research, a useful distinction can be drawn between quantitative urban and regional models. Quantitative urban models are concerned with internal city structure (the network of economic interactions within a single city). In contrast, quantitative regional models are instead concerned with systems of cities or regions (the network of economic interactions between cities or regions). The main distinction between these two classes of models is that their different spatial scales change the relative importance of different economic mechanisms. Between cities, goods trade and migration are typically the dominant mechanisms of interaction. Within cities, commuting (the separation of residence and workplace) becomes relatively more important.

Research on quantitative urban models includes Ahlfeldt et al. (2015), Dingel and Tintelnot (2020), Heblich et al. (2020), Owens III et al. (2020), Gechter and Tsivanidis (2022), Miyauchi et al. (2022), Kreindler and Miyauchi (2023), Monte et al. (2023), Nagy (2023), Severen (2023), Almagro and Domínguez-lino (2024), Redding and Sturm (2024), and Tsivanidis (2024), as surveyed in Redding and Rossi-Hansberg (2017), Redding (2023) and Redding (2025). Research on quantitative regional models includes Redding and Sturm (2008), Allen and Arkolakis (2014), Ahlfeldt et al. (2015), Redding (2016), Allen et al. (2024), Caliendo et al. (2018), Desmet et al. (2018), Monte et al. (2018), Fajgelbaum and Gaubert (2020), and Tsivanidis (2024), as reviewed in Redding and Rossi-Hansberg (2017), Allen and Arkolakis (2022a) and Allen and Arkolakis (2025).

Second, this article also contributes to the empirical literature on the impact of transport improvements on the location of economic activity, as reviewed in Redding and Turner (2015) and Donaldson (2025). One strand of this literature has used variation across cities and regions, including Chandra and Thompson (2000), Michaels (2008), Duranton and Turner (2011), Duranton and Turner (2012), Faber (2014), Duranton et al. (2014), Donaldson and Hornbeck (2016), Donaldson (2018), Baum-Snow et al. (2020), Hornbeck and Rotemberg (2024) and Weiwu (2024). A second group of studies have looked within cities, including Warner (1978), Jackson (1987), McDonald and Osuji (1995), Gibbons and Machin (2005), Baum-Snow and Kahn (2005), Baum-Snow (2007), Billings (2011), Baum-Snow et al. (2017), Brooks and Lutz (2018), Gonzalez-Navarro and Turner (2018), Baum-Snow (2020), and Heblich et al. (2020).

Most existing research evaluating transport improvements has focused on the positive economics question of the impact of a given transport improvement on the location of economic activity and welfare. More recent research has begun to compare alternative possible transport improvements and to evaluate the optimal transport network, including Fajgelbaum and Schaal

(2020), Allen and Arkolakis (2022b), Bordeu (2023), Kreindler et al. (2023), Almagro et al. (2024) and Jiang (2025). Other research on optimal spatial policy more broadly includes Fajgelbaum and Gaubert (2020), Fukui et al. (2024) and Mongey and Waugh (2024).

The remainder of the paper is structured as follows. Section 2 introduces our baseline quantitative urban model. Section 3 discusses the relationship with conventional cost-benefit analysis. Section 4 considers the relationship with market-access-based approaches. Section 5 illustrates the use of our baseline quantitative urban model to evaluate the impact of transport improvements using a numerical example. Section 6 summarizes our conclusions.

2 Quantitative Urban Model

We next outline a baseline quantitative urban model following Ahlfeldt et al. (2015). We consider a city that is embedded in a wider economy. The city consists of a discrete set of locations (city blocks) that are indexed by $n, i \in \mathbb{N}$. Time is discrete and is indexed by t. There are two types of agents: workers and landlords. Each worker is endowed with one unit of labor that is supplied inelastically. Workers have idiosyncratic preferences for each pair of residence and workplace and are geographically mobile within the city. Landlords are geographically immobile and own local floor space.

We consider two different assumptions about worker mobility with the wider economy: (i) A "closed-city" specification, with an exogenous measure of workers ($L_{\mathbb{N}t} = L_{\mathbb{N}}$), in which worker utility is endogenous; (ii) An "open-city" specification, in which the measure of workers ($L_{\mathbb{N}t}$) is endogenously determined by population mobility with a wider economy that provides a reservation level of utility $\overline{U}_t$. In this open-city specification, workers choose whether to move to the city before observing their idiosyncratic preferences. In both specifications, workers choose their preferred pair of residence and workplace within the city after observing their realizations for idiosyncratic preferences for each pair of residence and workplace. We assume a continuous measure of workers $L_{\mathbb{N}t}$, which ensures that the realized values of variables equals their expected values, and abstracts from any issues of granularity or small-sample variation.

Firms produce a single final good under conditions of perfect competition and constant returns to scale. This final good is costlessly traded and chosen as the numeraire. The model is static, but productivity, amenities, the supply of floor space and the transport network are allowed to evolve over time. To reduce notational clutter, we suppress the time subscript from now onwards, except where important.

2.1 Workplace-Residence Choices

Worker preferences are defined over the final consumption good and residential floor space. We assume that these preferences take the Cobb-Douglas form, such that the indirect utility for a worker ω residing in n and working in i is:

$$u_{ni}(\omega) = \frac{B_n b_{ni}(\omega) w_i}{\kappa_{ni} P_n^{\alpha} Q_n^{1-\alpha}}, \qquad 0 < \alpha < 1, \tag{1}$$

where $P_n = 1$ is the price of the final consumption good; Q_n is the price of residential floor space; w_i is the wage; κ_{ni} is an iceberg commuting cost; we model this iceberg commuting cost as depending on bilateral travel time (τ_{ni}) using the transport network: $\kappa_{ni} = e^{\kappa \tau_{ni}} \in [1, \infty)$, where $\kappa > 0$ parameterizes commuting costs;[1] B_n captures residential amenities, which can be endogenous to the surrounding concentration of economic activity through agglomeration forces; and $b_{ni}(\omega)$ is an idiosyncratic preference draw that captures all the idiosyncratic factors that can cause an individual to live and work in particular locations within the city.[2]

Residential amenities (B_n) are assumed to depend on residential fundamentals ($\overline{B}_n$) and residential externalities ($\mathbb{B}_n$). Residential fundamentals capture features of physical geography that make a location a more of less attractive place to live independently of the surrounding concentration of economic activity (e.g., green areas). Residential externalities capture the interactions between residents within the city (e.g., positive externalities through local public goods and negative externalities through crime):

$$B_n = \overline{B}_n \mathbb{B}_n^{\eta^B}, \qquad \mathbb{B}_n \equiv \sum_{i \in \mathbb{N}} e^{-\delta^B \tau_{ni}} R_i, \tag{2}$$

where R_i is the measure of residents in location i; η^B governs the magnitude of these residential externalities; and δ^B controls their spatial decay with travel time.

Idiosyncratic preferences ($b_{ni}(\omega)$) are drawn from an independent extreme value (Fréchet) distribution for each residence-workplace pair and for each worker:

$$G(b) = e^{-b^{-\epsilon}}, \qquad \epsilon > 1, \tag{3}$$

where we normalize the Fréchet scale parameter in equation (3) to one, because it enters the worker choice probabilities isomorphically to residential amenities (B_n) in equation (1); the

[1] Although commuting costs (κ_{ni}) are modelled in terms of utility, they enter the indirect utility function (1) multiplicatively with the wage, which implies that they are proportional to the opportunity cost of time. Therefore, similar results hold if commuting costs are instead modeled as a reduction in effective units of labor.

[2] A closely-related formulation assumes instead that workers have heterogeneous productivity for each pair of residence and workplace. Both specifications yield similar predictions for commuting probabilities, but imply different interpretations for observed wages. In the heterogeneous productivity specification, observed wages equal the wage per effective unit of labor multiplied by the average number of effective units of labor.

smaller the Fréchet shape parameter ϵ, the greater the heterogeneity in idiosyncratic preferences, and the less sensitive are worker location decisions to economic variables.[3]

These idiosyncratic preference shocks make solving the model's commuter market clearing condition tractable by ensuring that each residence-workplace pair faces an upward-sloping supply function for commuters in terms of wages adjusted for amenities, commuting costs and the cost of living $(B_n w_i / (\kappa_{ni} P_n^{\alpha} Q_n^{1-\alpha}))$. Using the properties of the extreme value distribution (3), the probability that a worker chooses to reside in n and work in i is:

$$\lambda_{ni} = \frac{L_{ni}}{L_{\mathbb{N}}} = \frac{(B_n w_i)^{\epsilon} \left(\kappa_{ni} P_n^{\alpha} Q_n^{1-\alpha}\right)^{-\epsilon}}{\sum_{k \in \mathbb{N}} \sum_{\ell \in \mathbb{N}} (B_k w_\ell)^{\epsilon} \left(\kappa_{k\ell} P_k^{\alpha} Q_k^{1-\alpha}\right)^{-\epsilon}}, \tag{4}$$

where L_{ni} is the measure of commuters from n to i; recall that $L_{\mathbb{N}}$ is the measure of workers in the city; and we report the derivations in Online Appendix B.1.

A key implication of equation (4) is that bilateral commuting flows satisfy a gravity equation, which is consistent with a large reduced-form empirical literature.[4] As in structural gravity models in international trade, bilateral commuting flows depend not only on "bilateral resistance" (κ_{ni}) between a pair of locations n and i in the numerator, but also on "multilateral resistance" ($\kappa_{k\ell}$ for all $k, \ell \in \mathbb{N}$) in the denominator. Although individual workers have idiosyncratic preferences for each residence-workplace pair, because there is a continuous measure of workers, there is no uncertainty in the supply of commuters for any residence-workplace pair.[5]

Summing across workplaces in equation (4), we obtain the probability of residing in each location ($\lambda_n^R = \sum_{\ell \in \mathbb{N}} \lambda_{n\ell}$):

$$\lambda_n^R = \frac{R_n}{L_{\mathbb{N}}} = \frac{(B_n)^{\epsilon} \, \Phi_n^R \left(P_n^{\alpha} Q_n^{1-\alpha}\right)^{-\epsilon}}{\sum_{k \in \mathbb{N}} (B_k)^{\epsilon} \, \Phi_k^R \left(P_k^{\alpha} Q_k^{1-\alpha}\right)^{-\epsilon}} \qquad \Phi_n^R \equiv \sum_{\ell \in \mathbb{N}} (w_\ell / \kappa_{n\ell})^{\epsilon}, \tag{5}$$

where recall that R_n is the measure of residents in location n; and we have defined Φ_n^R as a measure of residential commuting market access; which depends on commuting costs and the wages offered in each workplace.

Summing across residences in equation (4), we obtain the probability of being employed in each location ($\lambda_i^L = \sum_{k \in \mathbb{N}} \lambda_{ki}$):

$$\lambda_i^L = \frac{L_n}{L_{\mathbb{N}}} = \frac{(w_i)^{\epsilon} \, \Phi_i^L}{\sum_{\ell \in \mathbb{N}} (w_\ell)^{\epsilon} \, \Phi_\ell^L}, \qquad \Phi_i^L \equiv \sum_{k \in \mathbb{N}} B_k^{\epsilon} \left(\kappa_{ki} P_k^{\alpha} Q_k^{1-\alpha}\right)^{-\epsilon}, \tag{6}$$

[3] Modeling idiosyncratic preferences using the extreme value distribution has a long tradition in transportation economics, dating back to McFadden (1974). Here idiosyncratic preferences are specific to a residence-workplace pair, which allows for sorting by residence and workplace. But it is straightforward to consider a generalized extreme value specification with a nesting structure, in which for example workers choose a residence and then a workplace, with potentially different dispersion parameters for each decision (e.g., Baum Snow and Lu 2024).

[4] See for example McDonald and McMillen (2010) and Fortheringham and O'Kelly (1989).

[5] In the case of a discrete number of workers, "granularity" or small sample variation can become relevant for small spatial units, as analyzed in Dingel and Tintelnot (2020).

where L_i is the measure of workers employed in location i; and we have defined Φ_i^L as a measure of workplace commuting market access, which depends on the cost of living adjusted for amenities and commuting costs in each residence.

An additional implication of the extreme value distribution for idiosyncratic preferences (3) is that expected utility conditional on choosing a residence-workplace pair is the same across all residence and workplace pairs within the city:

$$U = \mathbb{E}\left[u\right] = \vartheta \left[\sum_{k \in \mathbb{N}} \sum_{\ell \in \mathbb{N}} \left(B_k w_\ell\right)^{\epsilon} \left(\kappa_{k\ell} P_k^{\alpha} Q_k^{1-\alpha}\right)^{-\epsilon} \right]^{\frac{1}{\epsilon}}, \tag{7}$$

where the expectation is taken over the distribution for idiosyncratic preferences; $\vartheta \equiv \Gamma((\epsilon - 1)/\epsilon)$; $\Gamma(\cdot)$ is the Gamma function; and the derivations are reported in Online Appendix B.1.

Therefore, each individual worker has a preferred residence-workplace pair based on their realizations for idiosyncratic preferences. But expected utility conditional on choosing a given residence-workplace pair is the same across all residence-workplace pairs. The intuition is as follows. On the one hand, higher amenity-adjusted real income $(B_i w_i / \left(\kappa_{ni} P_n^{\alpha} Q_n^{1-\alpha}\right))$ for residence n and workplace i raises the utility of a worker with a given realization for idiosyncratic preferences $b_{ni}\left(\omega\right)$, and hence increases expected utility. On the other hand, higher amenity-adjusted real income $(B_i w_i / \left(\kappa_{ni} P_n^{\alpha} Q_n^{1-\alpha}\right))$ attracts workers with lower realizations for idiosyncratic preferences $b_{ni}\left(\omega\right)$, which reduces expected utility. With a Fréchet distribution of idiosyncratic preferences, these two effects exactly offset one another. Residence-workplace pairs with high amenity-adjusted real income attract more commuters on the extensive margin until expected utility conditional on choosing a given residence-workplace pair is the same across all residence-workplace pairs.

2.2 Production

Markets are assumed to be perfectly competitive. The single final good can be costlessly traded both within the city and the wider economy. This final good is produced using labor and commercial floor space. We assume a constant returns to scale production technology that takes the Cobb-Douglas form. Cost minimization and zero-profits imply that price equals unit cost in each location with positive production:

$$1 = \frac{1}{A_n} w_n^{\beta} q_n^{1-\beta}, \qquad 0 < \beta < 1, \tag{8}$$

where A_n denotes productivity; q_n is the price of commercial floor space; and we have used our choice of numeraire ($P_n = 1$).

Productivity (A_n) is assumed to depend on production fundamentals ($\overline{A}_n$) and production externalities ($\mathbb{A}_n$). Production fundamentals capture features of physical geography that make a location a more or less attractive place to produce independently of the surrounding concentration of economic activity (e.g., access to natural water). Production externalities capture the interactions between workers within the city (e.g., knowledge spillovers):

$$A_n = \overline{A}_n \mathbb{A}_n^{\eta^A}, \qquad \mathbb{A}_n \equiv \sum_{i \in \mathbb{N}} e^{-\delta^A \tau_{ni}} L_i, \tag{9}$$

where η^A governs the magnitude of these production externalities and δ^A parameterizes their spatial decay with travel time.

2.3 Commuter Market Clearing

Commuter market clearing requires that the measure of workers employed in each workplace equals the measure of workers commuting to that workplace. From equations (4)-(6), we can write this commuter market clearing condition as:

$$L_i = \sum_{n \in \mathbb{N}} \lambda_{ni|n}^R R_n, \qquad \text{where} \qquad \lambda_{ni|n}^R = \frac{\lambda_{ni}}{\lambda_n^R} = \frac{(w_i/\kappa_{ni})^\epsilon}{\sum_{\ell \in \mathbb{N}} (w_\ell/\kappa_{n\ell})^\epsilon}. \tag{10}$$

Commuter market clearing also implies that average income per capita in each residence (v_n) is equal to the sum across workplaces of the wage in each workplace (w_i) multiplied by the probability of commuting to that workplace ($\lambda_{ni|n}^R$):

$$v_n = \sum_{i \in \mathbb{N}} \lambda_{ni|n}^R w_i. \tag{11}$$

2.4 Floor Space Market Clearing

We consider two alternative specifications of the market for floor space. First, we consider the case of a segmented market for floor space, in which the supplies of residential floor space (H_n^R) and commercial floor space (H_n^L) are both perfectly inelastic. Second, we examine the case of an integrated market for floor space, in which the overall supply of floor space (H_n) is perfectly inelastic, but floor space can be reallocated between residential use (H_n^R) and commercial use (H_n^L) to arbitrage away any differences in the return to these alternative uses. The overall supply of floor space (H_n) depends on both geographical land area and the ratio of floor space to geographical land area (as reflected in the height of buildings).[6]

Market clearing for residential floor space implies that income from the ownership of residential floor space equals payments for its use:

$$Q_n H_n^R = (1 - \alpha)\, v_n R_n. \tag{12}$$

[6] It is straightforward to allow the overall supply of floor space to be an increasing function of its price with a constant elasticity, as in Saiz (2010) and Heblich et al. (2020).

Similarly, market clearing for commercial floor space implies that income from the ownership of commercial floor space equals payments for its use:

$$q_n H_n^L = \frac{1-\beta}{\beta} w_n L_n, \tag{13}$$

where we used the property of the Cobb-Douglas production technology that payments to commercial floor space and labor are constant multiples of one another.

With a segmented market for floor space, the prices of residential and commercial floor space within a given location can differ from one another ($Q_n \neq q_n$). With an integrated market for floor space, the fractions of floor space allocated to residential use ($(1-\theta_n) = H_n^R / H_n$) and commercial use ($\theta_n = H_n^L / H_n$) are determined by the following no-arbitrage condition:

$$\begin{array}{llll} \theta_n & = 1 & \text{if} & q_n & > \xi_n Q_n \\ \theta_n & \in [0,1] & \text{if} & q_n & = \xi_n Q_n \\ \theta_n & = 0 & \text{if} & q_n & < \xi_n Q_n \end{array}, \tag{14}$$

where $\xi_n \geq 1$ is the tax equivalent of land use regulations.

We assume that floor space is owned by local landlords who are geographically immobile and consume only the traded final good. We abstract from idiosyncratic preference draws for landlords, because they are geographically immobile, and hence these preference draws would not affect the equilibrium allocation in any way.

2.5 General Equilibrium

We begin by considering general equilibrium with a closed-city and segmented markets for floor space. The equilibrium spatial distribution of economic activity within the city is determined by the model parameters (α, β, κ, ϵ, η_B, δ_B, η_A, δ_A) and the following exogenous location characteristics: residential fundamentals ($\overline{B}_n$), production fundamentals ($\overline{A}_n$), the supplies of residential and commercial floor space (H_n^R, H_n^L), and the transport network (τ_{ni}). Given these parameters and exogenous location characteristics, the closed-city general equilibrium is referenced by residents (R_n), employment (L_n), wages (w_n), average residential income (v_n), the prices of residential and commercial floor space (Q_n, q_n), and expected utility (U), given exogenous total city population ($L_{\mathbb{N}}$). From solutions for these equilibrium objects, all the other endogenous variables of the model can be determined.

We now show that the general equilibrium of the model admits a tractable theoretical characterization. In particular, the conditions for general equilibrium can be written in the form required to apply Theorem 1 from Allen et al. (2024) for uniqueness:

$$x_{nh} = \sum_{i \in \mathbb{N}} \mathcal{K}_{nih} \prod_{h' \in \mathbb{H}} x_{ih'}^{\gamma_{nhh'}}, \tag{15}$$

where x_{nh} is an endogenous variable; $\mathcal{K}_{nih}$ is a kernel that characterizes bilateral frictions; $n, i \in \mathbb{N}$ denote locations; and $h \in \mathbb{H}$ denote economic interactions, which here include residents, employment, and the prices of residential and commercial floor space. A sufficient condition for the existence of a unique equilibrium is that the spectral radius of a coefficient matrix of model parameters is less than or equal to one, as shown in Online Appendix B.2.

The determination of general equilibrium in the open-city with segmented markets for floor space is analogous, except that total city population ($L_{\mathbb{N}t}$) is endogenously determined by population mobility with the wider economy and its exogenous reservation level of utility ($\overline{U}$). The determination of general equilibrium with an integrated market for floor space is also analogous, except that there is an additional no-arbitrage condition between the prices of residential floor space (Q_n) and commercial floor space (q_n).

2.6 Counterfactuals

A key feature of quantitative urban models is that they are sufficiently rich so as to be able to rationalize the observed data as an equilibrium of the model. The model includes structural residuals that vary by location, such as production and residential fundamentals ($\overline{A}_n$, $\overline{B}_n$). These structural residuals can adjust by location, such that the observed data are consistent with the structural equations of the model.

This property of quantitative urban models typically implies that they are invertible, in the sense that given known parameters and the observed endogenous variables, we can back out unique values of the structural residuals such that the model is consistent with the observed data. Furthermore, this invertibility property can hold even in the presence of multiple equilibria, because these models condition on the observed equilibrium in the data. Given known model parameters, the observed endogenous variables and the equilibrium conditions of the model (e.g., cost minimization, zero profits and population mobility) can together contain enough information to uniquely determine these structural residuals, even though there could have been another equilibrium for the same value of the model parameters

Since quantitative urban models are able to rationalize the observed spatial distribution of economic activity as an equilibrium, they provide a suitable platform for undertaking counterfactuals for how realistic public-policy interventions would change this observed spatial distribution of economic activity. In our numerical example below, we consider a counterfactual for the construction of a railway network subway line that reduces bilateral commuting costs between some pairs of locations by more than for other pairs of locations.

In our baseline specification, we assume no agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$), such that productivity and amenities are determined solely by exogenous location characteristics. In this case, there exists a unique equilibrium, and our counterfactuals yield determinate predictions for the impact of the construction of the railway line on the spatial distribution of economic activity. We also consider an augmented specification with agglomeration forces, in which case productivity and amenities respond endogenously to changes in the spatial distribution of economic activity. In this case, for sufficiently strong agglomeration and dispersion forces, there is the potential for multiple equilibrium in the model. We solve for a counterfactual equilibrium using starting values for the model's endogenous variables from the initial equilibrium before the construction of the railway network.

In general, two main approaches have been used to undertaking counterfactuals in quantitative urban models. The first "covariates" approach estimates the model, recovers the predicted values of model variables such as commuting costs, and undertakes counterfactuals using these model predictions. The second "calibrated shares" approach undertakes counterfactuals conditioning on observed bilateral commuting flows in the initial equilibrium in the data. Following the international trade literature (Dekle et al. 2007 and Costinot and Rodríguez-Clare 2014), the calibrated shares approach is often referred to as "exact-hat algebra," because it rewrites the model's counterfactual equilibrium conditions in terms of the observed values of variables in the initial equilibrium in the data and the relative changes of variables ("hats") between the initial equilibrium and the counterfactual equilibrium. The approach is exact, in the sense that it solves for a counterfactual equilibrium using the full structure of the non-linear model, without any approximation. In our numerical example below, we know all location characteristics, such as commuting costs and production and residential fundamentals. Therefore, we directly solve for a counterfactual equilibrium using these known location characteristics.

We denote the value of a variable in the initial equilibrium without a prime (x), its value in the counterfactual equilibrium with a prime (x'), and the relative changes of variables between the two equilibria with a hat ($\widehat{x} = x'/x$). Given known location characteristics, we first solve for an initial equilibrium before a transport improvement. Given an assumed change in the transport network, and the resulting changes in commuting costs ($\widehat{\kappa}_{ni}$), we next solve for the counterfactual equilibrium after the transport improvement. In both cases, we solve the model's system of general equilibrium conditions, as discussed further in Online Appendix B.3.

In our numerical example below, we assume a closed-city and an integrated market for floor space. We assume positive and finite production and residential fundamentals. Additionally, in specifications with agglomeration forces, we assume positive spillovers of production and residential externalities. As a result, all locations have positive and finite productivity and amenities. Since the support of the Fréchet distribution for idiosyncratic preferences is unbounded from above, these assumptions ensure that all locations are incompletely specialized ($0 < \theta_n < 1$). Finally, we also assume no land use regulations ($\xi_n = 1$), which implies that the prices of residential and commercial floor space are equalized within each location ($Q_n = q_n$).

From equations (7) and (4), the increase in expected worker utility from the construction of the railway network can be expressed as:

$$\widehat{U} = \left[\sum_{k \in \mathbb{N}} \sum_{\ell \in \mathbb{N}} \lambda_{k\ell} \left(\widehat{B}_k \widehat{w}_\ell \right)^{\epsilon} \left(\widehat{\kappa}_{k\ell} \widehat{Q}_k^{1-\alpha} \right)^{-\epsilon} \right]^{\frac{1}{\epsilon}}, \tag{16}$$

where we used our choice of numeraire ($\widehat{P}_t = 1$).

The construction of the railway network also changes the value of land in each location. Since landlords are geographically immobile and the price of floor space differs across locations, this transport improvement has distributional consequences across the landlords in different locations. Since landlords consume only the numeraire final good, the change in the welfare of landlords in each location equals the change in the value of land.

3 Traditional Cost-Benefit Analysis

We now contrast the general equilibrium analysis of the impact of the transport improvement in our quantitative urban model in the previous section with conventional cost-benefit analyses that are typically undertaken in partial equilibrium.

The traditional approach to estimating the economic benefit from a transport improvement between a pair of locations n and i starts by estimating a Marshallian demand curve ($D(\tau, v)$) for travel between n and i given the travel cost (τ) and consumer income (v), as shown in Figure 1 below.[7] Given this demand curve and the estimated travel cost reduction from the transport improvement (($\tau_0 - \tau_1$)), the resulting welfare gain can be measured as the increase in consumer surplus (area (τ_0, a, b, τ_1)). This increase in consumer surplus includes both the substitution effect and the income effect from the lower travel cost given the consumer's income (v).

This increase in consumer surplus is closely related to conventional welfare concepts of compensating and equivalent variation. Compensating variation corresponds to the amount of income that would need to be taken away from the consumer after the transport improvement, in order for her to obtain the same level of utility as before the transport improvement. This corresponds to a change in the area under a Hicksian demand curve ($D(\tau, u)$) that holds utility rather than income constant (area (τ_0, a, d, τ_1) in Figure 2). Equivalent variation corresponds to the additional income that would need to be given to the consumer before the transport improvement, in order for her to obtain the same level of utility as after the transport improvement. Again this corresponds to a change in the area under a Hicksian demand curve ($D(\tau, u)$) that holds utility rather than income constant (area (τ_0, c, b, τ_1) in Figure 2). The gap between the two Hicksian demand curves in Figure 2 corresponds to the welfare gain from the transport improvement. As

[7]For a summary of this traditional approach, see for example Jones (1977).

Figure 1: Increase in Consumer Surplus from a Transport Improvement

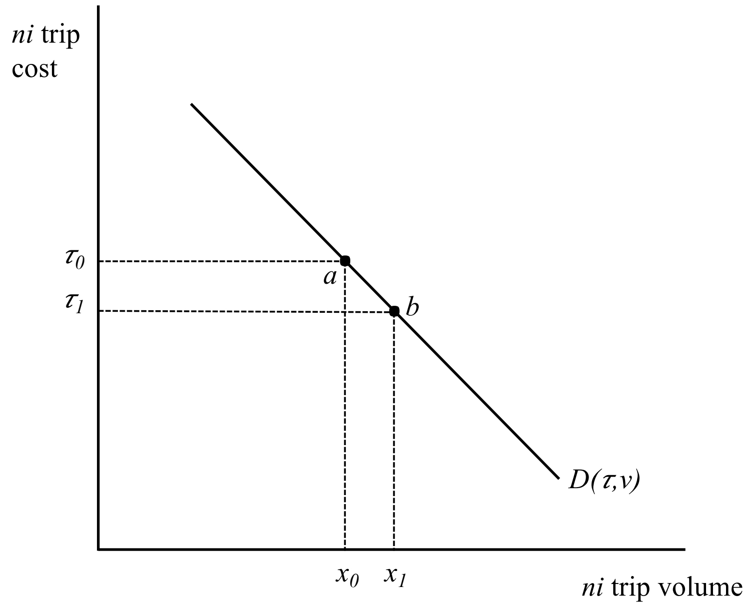

Note: τ_0 and x_0 are the travel cost and trip volume between locations n and i before the transport improvement, respectively; τ_1 and x_1 are the travel cost and trip volume between locations n and i after the transport improvement; the resulting increase in consumer surplus is area (τ_0, a, b, τ_1).

long as travel is a normal good (such that there is a positive income effect from the reduction in travel cost), the increase in consumer surplus from the transport improvement will lie in between the measures of compensating and equivalent variation, as shown in Figure 2.

Although conceptually straightforward, implementing this traditional cost-benefit approach in practice raises several challenges. In particular, one needs to evaluate the reduction in travel costs from the transport improvement, and estimate the demand curve for travel as a function of travel cost. Both are challenging if the transport improvement affects only one pair of locations, as in Figures 1 and 2. But, in reality, any given transport improvement will simultaneously affect many pairs of locations.

One special case that has received particular attention in the traditional cost-benefit literature is when (i) the demand for travel is perfectly inelastic and (ii) the change in travel cost equals the savings in travel time multiplied by the value of time. In this special case, the welfare gains from the transport improvement can be measured simply as the total value of the travel time saved. Although this measure is only exact when the demand for travel is perfectly inelastic, it provides a first-order approximation to the welfare gains more generally, since for small changes in travel time, the changes in the demand for travel are second-order small.[8]

When the changes to the transport network are large, such as the construction of an entire

[8]This special case is closely related to the social savings approach used by Fogel (1964) to evaluate the welfare gains from the U.S. railroad network in 1890.

Figure 2: Consumer Surplus, Compensating and Equivalent Variation

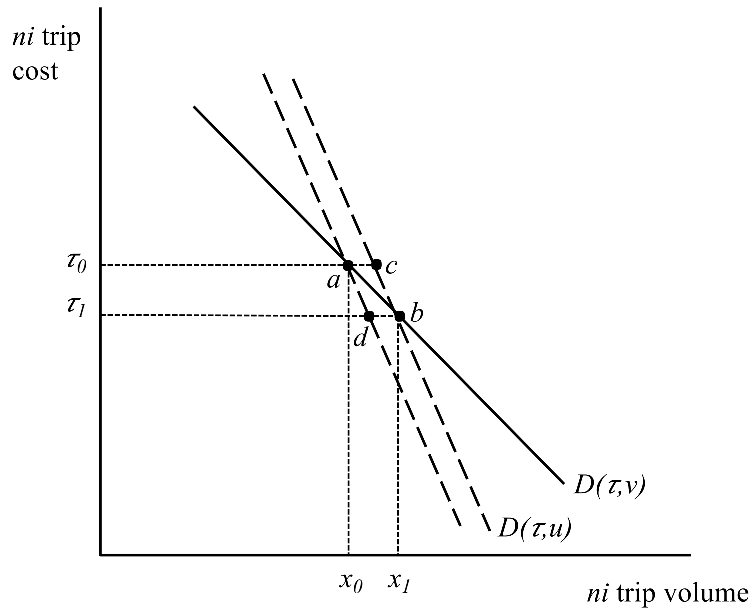

Note: Note: τ_0 and x_0 are the travel cost and trip volume between locations n and i before the transport improvement, respectively; τ_1 and x_1 are the travel cost and trip volume between locations n and i after the transport improvement; $D(\tau, y)$ is the Marshallian demand curve that holds income constant; $D(p, u)$ is the Hicksian demand curve that holds utility constant; area (τ_0, a, b, τ_1) is the increase in consumer surplus; area (τ_0, a, d, τ_1) is the compensating variation; and area (p_0, c, b, p_1) is the equivalent variation.

railway network, measuring the resulting welfare gains requires estimating the demand for travel between all pairs of locations affected by the transport improvement. A challenge in undertaking this estimation is that the demand for travel between any pair of locations is typically jointly determined in general equilibrium with the demand for travel between all other pairs of locations. As a result, the demand curve between any pair of locations n and i need not be stable, and can shift around based on changes in travel costs and the resulting redistribution of economic activity on other routes. Therefore, determining the demand for travel for each route ultimately involves solving for the spatial distribution of economic activity and determining travel on all routes in general equilibrium.

4 Market Access

Another alternative approach to solving for counterfactuals in a quantitative urban model involves the use of sufficient statistics based on measures of market access.

On the one hand, solving for counterfactuals in the quantitative urban model has several advantages. The researcher uses an internally-consistent equilibrium framework to evaluate the impact of counterfactual public policies. This approach makes explicit what assumptions are made, what is held constant, and what parameter values are used. Budget constraints and market clearing conditions necessarily hold. The solution for the counterfactual equilibrium is exact, because no approximation is made.

On the other hand, solving for counterfactuals in the quantitative urban model also has its disadvantages. Which predictions depend on the entire model structure versus which predictions hold in a broader class of models can be unclear. The sensitivity of counterfactual predictions to perturbations in model assumptions also can be unclear.

A sufficient statistics approach can address some of these limitations. Counterfactual predictions are derived from a smaller number of reduced-form equations that hold in a wider class of models. These counterfactual predictions only depend on the observed values of variables and assumptions about reduced-form parameters that in general are combinations of structural parameters in each model in that class.[9]

This class of quantitative urban models characterized by a gravity equation for bilateral commuting flows lends itself to a sufficient statistics representation in terms of the measures of residential market access (Φ_n^R) and workplace market access (Φ_n^L) introduced above, as shown in Tsivanidis (2024). Using the residential choice probability (5) and the workplace choice probability (6), we can rewrite residential and workplace market access (Φ_n^R, Φ_i^L) as follows:

$$\Phi_n^R = \frac{1}{\xi} \sum_{i \in \mathbb{N}} \kappa_{ni}^{-\epsilon} \frac{L_i}{\Phi_i^L}, \tag{17}$$

$$\Phi_i^L = \frac{1}{\xi} \sum_{n \in \mathbb{N}} \kappa_{ni}^{-\epsilon} \frac{R_n}{\Phi_n^R}, \tag{18}$$

as shown in Online Appendix B.2. Given data on employment (L_i) and residents (R_n), and estimates of bilateral commuting costs ($\kappa_{ni}^{-\epsilon}$), residential and workplace market access (Φ_n^R and Φ_i^L, respectively) can be recovered (up to scale) from this system of equations.

Using these relationships and the other general equilibrium conditions of the model, changes in employment ($\widehat{L}_n$), commercial floor prices ($\widehat{q}_n$), residents ($\widehat{R}_n$) and residential floor prices ($\widehat{Q}_n$) in response to the transport improvement can be written as log linear functions of changes in residential market access ($\widehat{\Phi}_n^R$), workplace market access ($\widehat{\Phi}_n^L$), and residuals:

$$\begin{aligned}
\log \widehat{L}_n &= \frac{1}{1+\epsilon\,(1-\beta)} \log \widehat{\Phi}_n^L + \log \widehat{e}_n^L, \\
\log \widehat{q}_n &= \frac{\beta}{1+\epsilon\,(1-\beta)} \log \widehat{\Phi}_n^L + \log \widehat{e}_n^q, \\
\log \widehat{R}_n &\approx \frac{\alpha}{1+\epsilon\,(1-\alpha)} \log \widehat{\Phi}_n^R + \log \widehat{e}_n^R, \\
\log \widehat{Q}_n &\approx \frac{1+\epsilon}{\epsilon\,(1+\epsilon\,(1-\alpha))} \log \Phi_n^R + \log \widehat{e}_n^Q,
\end{aligned} \tag{19}$$

[9]For a review of the public finance literature on sufficient statistics, see Chetty (2009).

where recall that a hat above a variable denotes a relative change between the counterfactual and initial equilibria ($\widehat{x} = x'/x$); and the derivations are reported in Online Appendix C.

Therefore, we can use equations (17) and (18) and the system of equations (19) to solve for the changes in residential market access ($\widehat{\Phi}_n^R$), workplace market access ($\widehat{\Phi}_i^L$), employment ($\widehat{L}_n$), commercial floor prices ($\widehat{q}_n$), residents ($\widehat{R}_n$) and residential floor prices ($\widehat{Q}_n$) induced by a transport improvement that changes commuting costs ($\widehat{\kappa}_{ni}$).

The first two relationships for employment ($\widehat{L}_n$) and commercial floor prices ($\widehat{q}_n$) in equation (19) are exact. The second two relationships for residents ($\widehat{R}_n$) and residential floor prices ($\widehat{Q}_n$) involve an approximation around an equilibrium with prohibitive commuting costs (in which $\widehat{v}_n \approx w_n \approx \left(\widehat{\Phi}_n^R\right)^{1/\epsilon}$).[10] The residuals ($\widehat{e}_n^L$, $\widehat{e}_n^q$, $\widehat{e}_n^R$, $\widehat{e}_n^Q$) depend on changes in productivity ($\widehat{A}_n$), amenities ($\widehat{B}_n$), and the supplies of commercial ($\widehat{H}_n^L$) and residential ($\widehat{H}_n^R$) floor space as follows:

$$\begin{aligned}
\log \widehat{e}_n^L &= \kappa^L + \frac{\epsilon}{1+\epsilon(1-\beta)} \log \widehat{A}_n + \frac{\epsilon(1-\beta)}{1+\epsilon(1-\beta)} \log \widehat{H}_n^L, \\
\log \widehat{e}_n^q &= \kappa^q + \frac{1+\epsilon}{1+\epsilon(1-\beta)} \log \widehat{A}_n - \frac{\beta}{1+\epsilon(1-\beta)} \log \widehat{H}_n^L, \\
\log \widehat{e}_n^R &= \kappa^R + \frac{\epsilon}{1+\epsilon(1-\alpha)} \log \widehat{B}_n + \frac{\epsilon(\alpha-1)}{1+\epsilon(1-\alpha)} \log \widehat{H}_n^R, \\
\log \widehat{e}_n^Q &= \kappa^Q + \frac{\epsilon}{1+\epsilon(1-\alpha)} \log \widehat{B}_n + -\frac{1+2\epsilon(\alpha-1)}{1+\epsilon(1-\alpha)} \log \widehat{H}_n^R,
\end{aligned} \tag{20}$$

where (κ^L, κ^q, κ^R, κ^Q) are constants that depend on the change in the common level of expected utility across all locations ($\widehat{U}$).

Two sets of implications follow from this system of equations (19). First, assume no agglomeration forces ($\widehat{A}_n = 1$ and $\widehat{B}_n = 1$) and a segmented market for floor space with exogenous supplies of commercial and residential floor space ($\widehat{H}_n^L = 1$ and $\widehat{H}_n^R = 1$). Under these assumptions, the changes in residential and workplace market access ($\widehat{\Phi}_n^R$, $\widehat{\Phi}_n^L$) are sufficient statistics for the reorganization of economic activity in response to transport improvements, up to the quality of the approximation around an equilibrium with no commuting costs. Therefore, under these assumptions, one can predict the impact of a transport improvement by solving for changes in market access, without necessarily having to solve the full system of general equilibrium conditions for a counterfactual equilibrium.

Second, assume agglomeration forces ($\widehat{A}_n \neq 1$ or $\widehat{B}_n \neq 1$) and/or integrated floor space markets with endogenous allocations of floor space between commercial and residential use ($\widehat{H}_n^L \neq 1$ and $\widehat{H}_n^R \neq 1$). Under these assumptions, the reorganization of economic activity in response to a transport improvement is determined by changes in both residential and workplace market access ($\widehat{\Phi}_n^R$, $\widehat{\Phi}_n^L$) and the residuals ($\widehat{e}_n^L$, $\widehat{e}_n^q$, $\widehat{e}_n^R$, $\widehat{e}_n^Q$). The changes in these residuals depend on changes

[10] Similar relationships hold in a version of the quantitative urban model in Section 2 in which the idiosyncratic shock ($b_{ni}(\omega)$) in equation (1) is to worker productivity rather than preferences, as shown in Tsivanidis (2024).

in production and residential externalities and endogenous reallocations of floor space between alternative uses. In general, determining the changes in these residuals typically requires solving for counterfactuals in the full non-linear model.[11]

Therefore, the extent to which residence (Φ_n^R) and workplace (Φ_n^L) market access are sufficient statistics for internal city structure depends on assumptions about agglomeration economies and the supplies of commercial and residential floor space. In the next section, we examine the quantitative relevance of the differences between the predictions from model counterfactuals and measures of market access for the impact of a transport improvement.

5 Quantitative Illustration

We now illustrate the use of our quantitative urban model to evaluate the impact of a transport improvement using a numerical example of city. By focusing on this numerical example, we consider a setting in which we know the true data generating process (DGP) and model parameters. Therefore, the data are generated according to the model, and we can examine the success of alternative approaches in approximating the true impact of the transport improvement. Our numerical example is motivated by the analysis of the construction of London's 19th-century railway network in Heblich et al. (2020). We calibrate some model parameters using empirical moments from that historical setting.

In Section 5.1, we introduce the structure of the city. In Section 5.2, we discuss the parameterization of the model. In Section 5.3, we examine the spatial distribution of economic activity in the initial equilibrium before the construction of the railway network. In Section 5.4, we undertake a counterfactual for the construction of the railway network. In Section 5.5, we examine the impact of this new transport technology on the distribution of commuting distances and travel times. In Section 5.6, we compare the counterfactual predictions from the solution of the full non-linear model with those based on changes in market access. In Section 5.7, we examine the aggregate implications of the new transport technology by computing changes in expected worker utility and aggregate land values.

[11] If production and residential externalities forces depend only on employment and residents in the *own* location, and the supplies of commercial and residential floor space are exogenous, changes in production and residential externalities can be expressed solely in terms of changes in residence and workplace market access, up to the first-order approximation used above. More generally, when production and residential externalities spill over across locations, this is no longer the case, because the spatial interactions between locations also depend on the decay of production and residential externalities.

5.1 City Structure

We consider a numerical example of a city that consists of a set of grid points arrayed in (x, y) space. We assume a grid of 22×22 points that are each one unit of distance apart from one another. We interpret one unit of distance in the model as one kilometer (km), such that the internal area bounded by this grid is $21 \times 21 = 441$ km^2, which is somewhat larger than the area of the County of London of around 314 square km^2.

We assume exogenous differences in production fundamentals ($\overline{A}_i$) that promote the concentration of economic activity in a central location, consistent with many real-world cities forming around natural advantages, such as ports or navigable rivers. We assume that the central node (11,11) has production fundamentals of $\overline{A}_i = 4$; surrounding nodes (9:13,9:13) have production fundamentals of $\overline{A}_i = 2$; and all other locations have production fundamentals of $\overline{A}_i = 1$. We assume that all locations have the same residential fundamentals of $\overline{B}_i = 1$. In our baseline specification, we abstract from agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$). In an augmented specification, we incorporate agglomeration forces, in which case the exogenous differences in production fundamentals induce endogenous differences in productivity and amenities (A_i and B_i) through production and residential externalities ($\mathbb{A}_i$ and $\mathbb{B}_i$). We assume that locations have the same area ($K_i = 100$), and the same density of development (ratio of floor space to area) of $\varphi_i = 1$, such that the exogenous supply of floor space in all locations is $H_i = \varphi_i K_i = 100$.

The transport network is modelled as the set of arcs that connect the grid points (nodes). Each arc has a weight that corresponds to the cost of travelling on that arc. We compute travel times between each bilateral pair of grid points as the least-cost paths across the arcs connecting those grid points. In the initial equilibrium, walking is the only mode of transport. All grid points are connected by this mode of transport, with a travel cost of 1 per unit distance (km). We consider a counterfactual for the construction of a railway network that consists of a central vertical line, a central horizontal line, a diagonal line, and an inverse diagonal line. Therefore, the railway network has a hub and spoke structure around the city center, as often observed for many real-world transport networks. We assume that the cost of travelling along an arc connected to the railway network is $1/\gamma$ per unit distance, where $\gamma > 1$.

5.2 Parameterization

We calibrate the model's parameters based on estimates from the existing empirical literature and empirical moments from the construction of London's 19th-century railway network from Heblich et al. (2020). We set the share of consumer expenditure on residential floor space $(1 - \beta)$ equal to 0.25, which is consistent with the estimates in Davis and Ortalo-Magné (2011). We assume that the share of firm expenditure on commercial floor space $(1 - \alpha)$ is 0.20, which is in line with the findings of Valentinyi and Herrendorf (2008). We set the dispersion of idiosyncratic preferences (ϵ) equal to 5, which is close to the estimate of 5.25 in Heblich et al. (2020), and lies in the center of the range of values from 2.18 to 8.3 in Ahlfeldt et al. (2015), Dingel and Tintelnot (2020), Severen (2023) and Kreindler and Miyauchi (2023).

We set the commuting cost semi-elasticity (κ) equal to 0.20, which together with our parameter choices ensures that the model matches the empirical finding in Heblich et al. (2020) that more than 90 percent of workers lived within 5 km of their workplace before the railway age.[12] We set the reduction in travel cost from a railway connection (γ) equal to 5, which together with our parameter choices implies that the model matches the empirical finding in Heblich et al. (2020) that around 50 percent of workers lived within 5 km of their workplace by the end of the construction of London's 19th-century railway network. We interpret this travel cost as the relative travel time of the two modes of transport, where our calibrated value of 5 is close to the relative travel time of walking and railways of 7 used in Heblich et al. (2020).

In our baseline specification, we assume no agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$). In our augmented specification with agglomeration forces, we assume elasticities of productivity and amenities with respect to travel-time-weighted employment and residential density of $\eta_A = \eta_B = 0.10$, and exponential rates of decay of production and residential externalities of $\delta_A = \delta_B = 0.08$. These assumed elasticities (η_A, η_B) are marginally higher than the range of 3-8 percent reported for production externalities in Rosenthal and Strange (2004) based on cross-city variation. But the common assumed elasticity of $\eta_A = \eta_B = 0.10$ is lower than the estimated residential elasticity of $\eta_B = 0.15$ in Ahlfeldt et al. (2015), and lower than the estimated elasticities in several studies using quasi-experimental sources of variation, including Greenstone et al. (2010) and Kline and Moretti (2014).[13] The assumed exponential rates of decay of $\delta_A = \delta_B = 0.08$ are smaller in absolute value than those estimated by Ahlfeldt et al. (2015), implying less localized production and residential externalities.

We show that our main quantitative findings for the role of transport improvements in allowing locations to specialize as workplaces and residences are robust across these specifications with and without agglomeration forces.

[12] The implied semi-elasticity of commuting flows to distance is $\epsilon\kappa = 0.20 \times 5$, which is somewhat large compared to empirical estimates using data measured in kilometers, but the interpretation of this semi-elasticity in the model depends on the interpretation of the units in which distance is measured.

[13] For reviews of the range of estimated agglomeration elasticities in existing empirical research, see the meta-analyses in Melo et al. (2009) and Ahlfeldt and Pietrostefani (2019).

Table 1: Parameterization

	Parameter	Value	Source
Without Agglomeration Forces			
Residential floor space expenditure share	$1-\alpha$	0.25	Ahlfeldt et al. (2015)
Commercial floor space cost share	$1-\beta$	0.20	Ahlfeldt et al. (2015)
Dispersion idiosyncratic preferences	ϵ	5	Heblich et al. (2020)
Commuting cost semi-elasticity	κ	0.20	Share commute < 5 km before rail
Railway reduction commuting cost	γ	5	Share commute < 5 km after rail
With Agglomeration Forces			
Production agglomeration forces	η_A	0.10	< Ahlfeldt et al. (2015)
Production agglomeration decay	δ_A	0.08	< Ahlfeldt et al. (2015)
Residential agglomeration forces	η_B	0.10	< Ahlfeldt et al. (2015)
Residential agglomeration decay	δ_B	0.08	< Ahlfeldt et al. (2015)

Note: Calibrated model parameters and the source for each calibrated parameter value; the specification without agglomeration forces uses $\eta_A = \eta_B = \delta_A = \delta_B = 0$; the specification with agglomeration forces uses the parameter values listed in the bottom panel of the table.

5.3 Initial Equilibrium

Figure 3 shows the initial equilibrium distribution of economic activity before the construction of the railway network for our baseline specification without agglomeration forces.[14] We indicate levels of economic activity in each location using a heatmap, in which higher values are denoted with lighter colors (more yellow), and lower values are denoted with darker colors (more blue).

As shown in the top-left panel, we find the gradient of the price of floor space ($\log Q_i = \log q_i$) has an approximately monocentric structure. There is a central peak (shown in yellow) at the location with the highest production fundamentals (11,11). Around this central peak, there is a square area of lower prices of floor space (shown in teal) at the inner-city locations with intermediate production fundamentals (9:13, 9:13). Beyond the boundaries of these inner-city locations with intermediate production fundamentals, there is a continuing gradient in the price of floor space (from light blue to dark blue), which reflects the impact of commuting costs. Locations closer to the center have lower commuting costs to employment concentrations in the center of the city. In order for expected utility to be equalized across all locations, these lower commuting costs must be compensated in equilibrium by higher prices of floor space.

In the top-right panel, we show the ratio of workers to residents ($\log (L_i/R_i)$), which captures locations' patterns of specialization as workplaces and residences. Locations with values of this ratio greater than one are net importers of commuters, while locations with values of this ratio less than one are net exporters of commuters. In the remaining two panels of the figure, we show the two separate components of this ratio. The bottom-left shows workers ($\log L_i$) from

[14] In Online Appendix D, we show the analog of Figure 3 and all subsequent figures for our augmented specification with agglomeration forces, and demonstrate a similar pattern of results.

the numerator. The bottom-right panel shows residents ($\log R_i$) from the denominator.

Figure 3: Initial Equilibrium Before the Railway Network

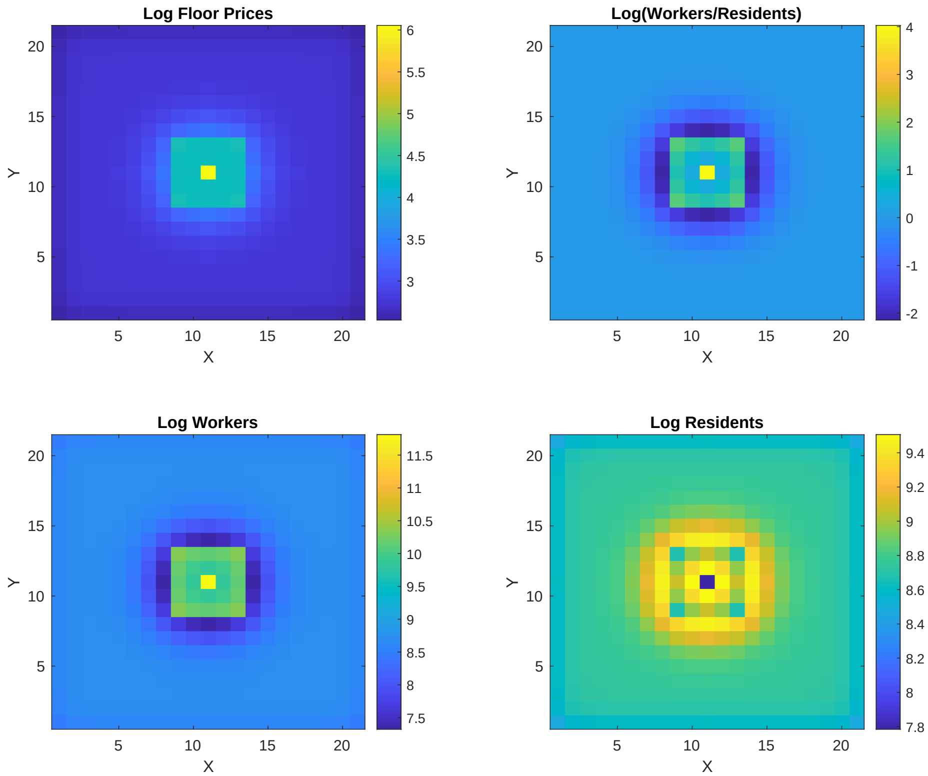

Note: Figure shows the initial equilibrium before the railway network in our baseline specification without agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$). City composed of 21×21 grid points one unit apart (interpreted as 1 kilometer). Central node (11,11) has production fundamentals (a_i) equal to 4. Surrounding nodes (9:13,9:13) have $a_i = 2$. All other locations have $a_i = 1$. All locations have residential fundamentals (b_i) equal to 1. All locations have area (K_i) equal to 100, a density of development (φ_i) equal to 1, and a supply of floor space equal to $H_i = \varphi_i K_i = 100$. In the initial equilibrium, walking is the only mode of transport, with a travel cost of 1 per unit distance. Top-left panel shows the log price of floor space ($\log Q_i = \log q_i$). Top-right panel shows the log ratio of workers to residents ($\log (L_i/R_i)$). Bottom-left panel shows log residents ($\log R_i$). Bottom-right panel shows log workers ($\log L_i$). Negative values are possible, because all variables are logged.

We find that the central location has the highest ratio of workers to residents (top right), the largest concentration of workers (bottom left), and the smallest concentration of residents (bottom right). This pattern reflects the interaction between the comparative advantage of locations as workplaces or residences and commuting costs. The central location as the highest production fundamentals relative to residential fundamentals, and hence specializes as a workplace, importing commuters from other surrounding locations. As employment concentrates in this central location to take advantage of its high productivity, this bids up the price of floor space, reducing the concentration of residents in that location. Nevertheless, for positive and finite levels of production and residential fundamentals, the central location remains incompletely specialized with positive values of both employment and residents, because the support of the Fréchet distribution for idiosyncratic preferences is unbounded from above. Therefore, there is a positive measure of workers who draw sufficiently high idiosyncratic preferences that they are willing to live in the central location despite its relatively high price of floor space.

Around the central location, we find a non-monotonic pattern of concentric rings of specialization, in which intermediate locations just beyond the boundaries of the inner city have the lowest ratios of workers to residents (top right), the smallest concentrations of workers (bottom left), and the largest concentrations of residents (bottom right). Again this pattern reflects the interaction between comparative advantage and commuting costs. These intermediate locations have lower production fundamentals than central locations, the same residential fundamentals as all locations, and lower commuting costs to the center than outlying locations. Therefore, these intermediate locations specialize as residences rather than workplaces, and are the largest sources of commuters for the central locations that specialize as workplaces. Nevertheless, all locations remain incompletely specialized with positive values of employment and residents, because the distribution of idiosyncratic preferences has a support that is unbounded from above.

5.4 Counterfactual Equilibrium

Figure 4 shows the counterfactual equilibrium distribution of economic activity after the construction of the railway network, again for our baseline specification without agglomeration forces. As in the previous figure, we indicate levels of economic activity in each location using a heatmap, in which higher values are denoted with lighter colors (more yellow), and lower values are denoted with darker colors (more blue).

We find that the construction of the railway network increases the price of floor space in the central location with the highest production fundamentals (compare the upper limit of the scale in the top-left panels of Figures 3 and 4). As the railway network reduces commuting costs, this allows the central location to further exploit its comparative advantage as a workplace rather than a residence, and import commuters from other locations. As this specialization occurs, and floor space is reallocated towards a more efficient use, this bids up the price of floor space in the central location.

The railway network also increases floor prices in outlying locations that are close to railway lines (with the structure of the railway network reflected in a "union jack" pattern of light blue areas for the price of floor space in the top-left panel of Figure 4). These outlying locations close to railway lines now have lower commuting costs to employment concentrations in the center of the city than other outlying locations. These lower commuting costs attract residents and bid up the price of floor space close to railway lines, until expected utility is again equalized across all locations. In contrast, outlying locations furthest from railway lines continue to have similar prices of floor space as before the construction of the railway network (comparing the lower limit of the scale in the top-left panels of Figures 3 and 4). As a result, the new transport technology increases the inequality in the price of floor space across locations.

Figure 4: Counterfactual Equilibrium after the Railway Network

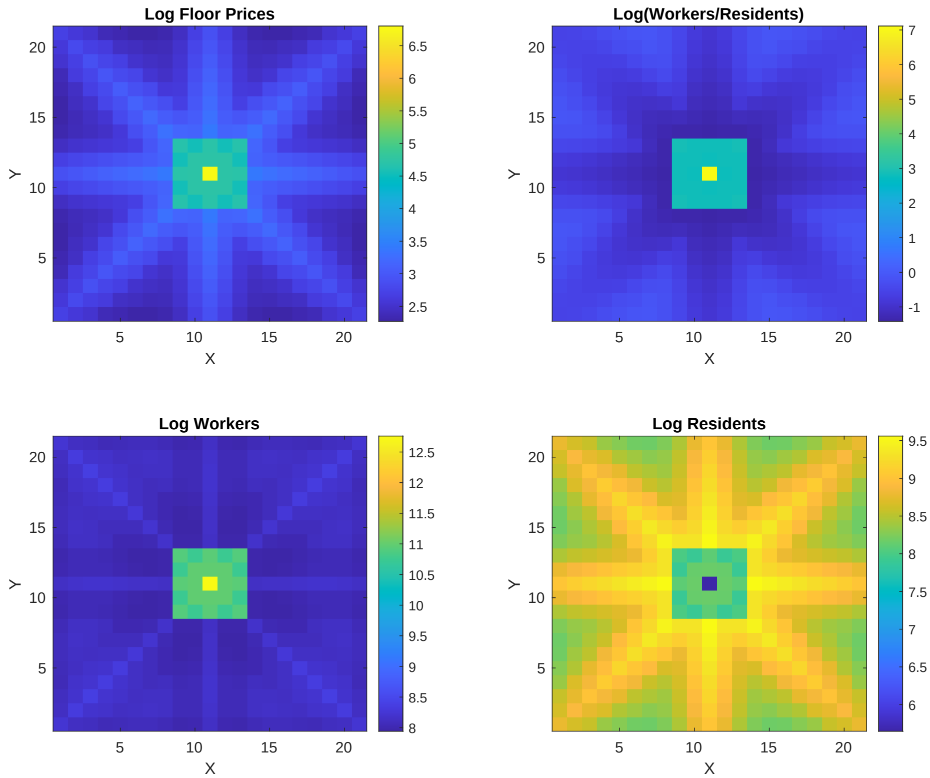

Note: Figure shows the counterfactual equilibrium after the railway network in our baseline specification without agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$). City composed of 21×21 grid points one unit apart (interpreted as 1 kilometer). Central node (11,11) has production fundamentals (a_i) equal to 4. Surrounding nodes (9:13,9:13) have $a_i = 2$. All other locations have $a_i = 1$. All locations have residential fundamentals (b_i) equal to 1. All locations have area (K_i) equal to 100, a density of development (φ_i) equal to 1, and a supply of floor space equal to $H_i = \varphi_i K_i$. In the counterfactual, vertical horizontal, diagonal and inverse diagonal railway lines are constructed (such that the railway network forms a "union jack." Railway lines reduce the cost of travel from 1 to $1/\gamma$ per unit distance (where $\gamma > 1$). Top-left panel shows the log price of floor space ($\log Q_i = \log q_i$). Top-right panel shows the log ratio of workers to residents ($\log (L_i/R_i)$). Bottom-left panel shows log residents ($\log R_i$). Bottom-right panel shows log workers ($\log L_i$). Negative values are possible, because all variables are logged.

The role of the railway network in enabling increased specialization to take advantage of patterns of comparative advantage is also evident in Figure 4 from the ratio of workers to residents (top right), the spatial distribution of workers (bottom left), and the spatial distribution of residents (bottom right). The construction of the railway network increases the ratio of workers to residents in the most central locations, as they increasingly specialize as workplaces (upper limit of the scale in the top-right panel). The concentric ring pattern of specialization in intermediate locations around the inner city in Figure 3 is still evident to some degree in Figure 4. These intermediate locations surrounding the areas with higher production fundamentals in the inner city (the areas surrounding locations 9:13, 9:13) have some of the highest concentrations of residents (bottom right) and lowest ratios of workers to residents (top right).

But this concentric ring pattern of specialization is now supplemented with a radial pattern of specialization in Figure 4. Outlying locations close to railway lines now have some of the largest concentrations of residents (bottom right) and some of the lowest ratios of workers to residents (top right). This concentration of residents along railway lines (bottom right) is also reflected to a more limited extent in the concentration of employment along these lines (bottom left). This pattern reflects the gravity structure of commuting flows. Firms located close to railway lines have large supplies of nearby residents living close to the railway lines, which in the presence of commuting costs reduces the wages that they need to pay to attract workers, and hence increases employment in these locations. Whereas outlying areas are relatively undifferentiated in Figure 3 before the construction of the railway network, they are now substantially more heterogeneous in Figure 4 after its construction. Some outlying locations furthest from railway lines experience declines in both workers and residents following the construction of the railway network, as they become relatively less attractive as workplaces and residences compared to other locations (comparing the lower limit of the scales in the bottom two panels in Figures 3 and 4).

5.5 Commuting Distributions

We now examine these implications of these changes in locations' patterns of specialization as workplaces and residences for commuting patterns.

Figure 5 shows the cumulative distribution of commuting distances in the initial equilibrium (blue) and counterfactual equilibrium (red), again for our baseline specification without agglomeration forces. We construct this cumulative distribution as follows. For each residence-workplace pair, we observe the equilibrium number of commuters and the bilateral distance travelled. We sort residence-workplace pairs by the distance travelled, and compute the cumulative sum of commuters for each distance travelled, divided by the total number of commuters in the city, which yields the share of workers who commute less than each distance.

As discussed above, we calibrated the model's parameters to match empirical moments from the construction of London's 19th-century railway network in Heblich et al. (2020). Before the construction of the railway network, more than 90 percent of workers live within 5km of their workplace in Figure 5. After the construction of the railway network, around 50 percent of workers live more than 5km from their workplace in Figure 5. Therefore, the reduction in commuting costs from the construction of the railway network leads to a substantial increase in the fraction of workers that commute over longer distances.

Figure 6 shows the cumulative distribution of effective distances, where effective distance adjusts for the different travel costs of walking and the railway, and has an interpretation as travel time. We construct this cumulative distribution in the same way as for the previous figure, but use effective distance (travel times) instead of distance. We show this cumulative distribution using

Figure 5: Cumulative Commuting Distance Distributions

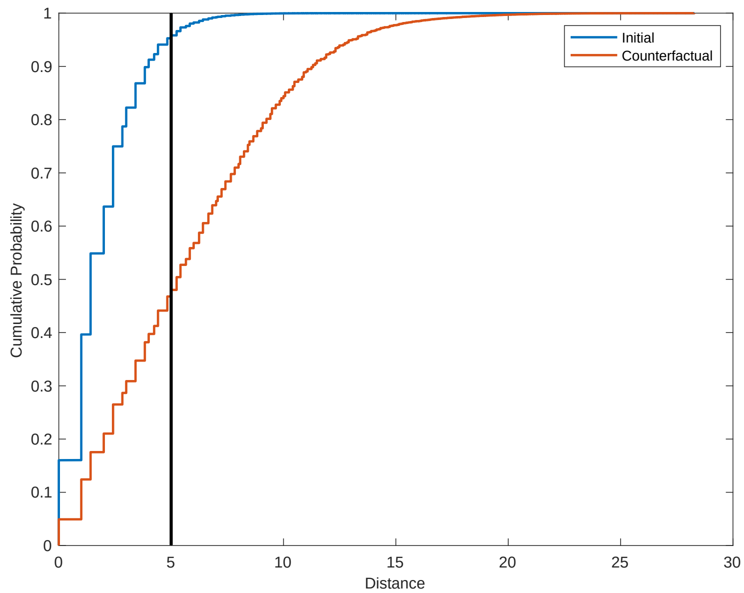

Note: Figure shows the cumulative distribution for the share of workers who commute less than each distance in the initial equilibrium (blue) and the counterfactual equilibrium (red) in the specification without agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$). Vertical black line shows 5 kilometers distance. In the initial equilibrium, walking is the only mode of transport, with a travel cost of 1 per unit distance. In the counterfactual, vertical horizontal, diagonal and inverse diagonal railway lines are constructed (such that the railway network forms a "union jack." Railway lines reduce the cost of travel from 1 to $1/\gamma$ per unit distance (where $\gamma > 1$).

initial commuting probabilities and initial travel times (blue), using initial commuting probabilities and counterfactual travel times (orange dashed line), and using counterfactual commuting probabilities and counterfactual travel times (red line).

The direct effect of the construction of the railway network is the reduce the amount of time taken to travel a given distance. Therefore, when we use the initial commuting probabilities for both the initial and counterfactual travel times (comparing the blue to the orange dashed line), we find a substantial fall in the share of workers who commute less than each travel time, as reflected in a leftwards shift in the cumulative distribution. However, our quantitative urban model predicts that workers and firms respond to the construction of the railway network by adjusting their location choices, with the reduction in commuting costs leading workers to increasingly separate their residence and workplace and to commute over longer distances. As a result, when we compare the initial distribution (using initial values for both commuting probabilities and travel times (blue line)) to the counterfactual distribution (using counterfactual values for both commuting probabilities and travel times (red line)), we find relatively little change in the share of workers who commute less than each travel time.

The extent to which workers respond to longer commuting costs by increasing commuting distance depends on model parameters, such as the dispersion in idiosyncratic preferences (ϵ) and the semi-elasticity of commuting costs with respect to travel times (κ). Nevertheless, this

Figure 6: Cumulative Commuting Effective Distance (Travel Time) Distributions

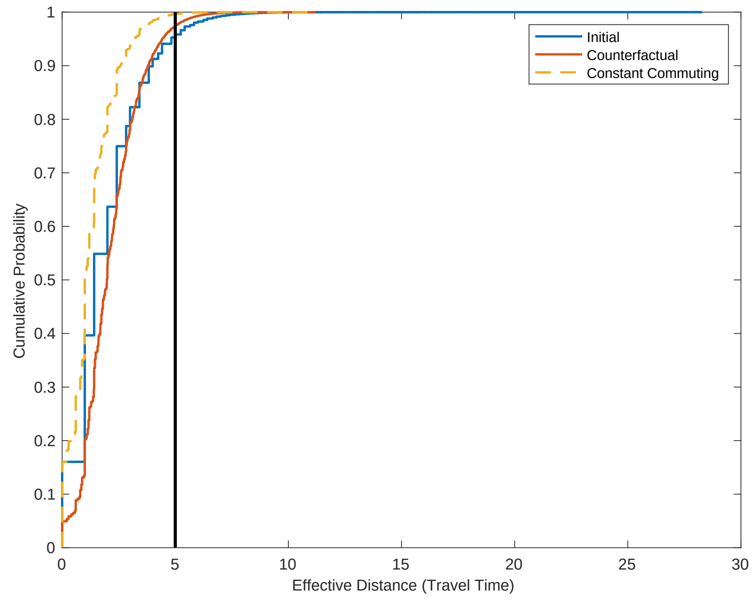

Note: Figure shows the cumulative distribution for the share of workers who commute less than each effective distance (travel time) in our baseline specification without agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$); blue line shows this cumulative distribution in the initial equilibrium (using initial commuting probabilities and travel times); orange dashed line shows this distribution using initial commuting probabilities and counterfactual travel times; red line shows this cumulative distribution in the counterfactual equilibrium (using counterfactual commuting probabilities and travel times). Vertical black line shows 5 kilometers distance. In the initial equilibrium, walking is the only mode of transport, with a travel cost of 1 per unit distance. In the counterfactual, vertical horizontal, diagonal and inverse diagonal railway lines are constructed (such that the railway network forms a "union jack"). Railway lines reduce the cost of travel from 1 to $1/\gamma$ per unit distance (where $\gamma > 1$). Effective distance adjusts for γ and has an interpretation as travel time.

pattern of results highlights that transport improvements need not reduce commuting times in equilibrium, because of an endogenous reorganization of economic activity in response to the transport improvement.[15] As a result, the direct impact of the new transport technology on travel times (based on initial commuting shares) can be quite misleading for its general equilibrium impact (based on counterfactual commuting shares), once this reorganization is taken into account. Quantitative urban models provide the theoretical structure required to solve for the general equilibrium reorganization of economic activity in response to the transport improvement. In contrast, it is challenging to capture this reorganization in conventional cost-benefit analyses, which are typically undertaken in partial equilibrium.

5.6 Comparison with Market Access Predictions

We next examine the ability of predictions based on market access to capture this general equilibrium reorganization of economic activity in response to the transport improvement.

[15]This point is related to the "fundamental law" of highway congestion, whereby increasing the capacity of a road leads to a proportionate increase in traffic, as examined empirically in Duranton and Turner (2011).

Figure [7] shows log relative changes between the counterfactual and initial equilibria ($\log \widehat{x}_n$) for employment (top left), commercial floor space prices (top right), residents (bottom left), and residential floor space prices (bottom right). Again we report results here for our baseline specification without agglomeration forces, while Online Appendix D demonstrates a similar pattern of results for our augmented specification incorporating agglomeration forces.

Figure 7: Counterfactual Versus Market Access Predictions

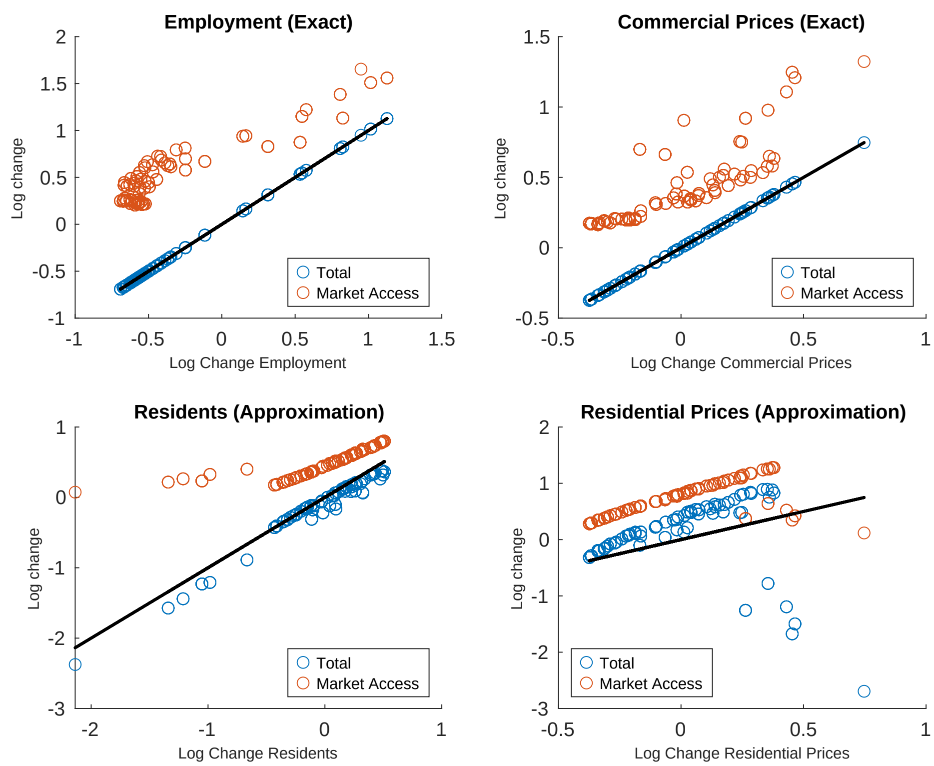

Note: Figure shows predicted log changes in each variable between the counterfactual and initial equilibria from model counterfactuals (horizontal axis) versus market access predictions (vertical axis) in the specification without agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$). Blue circles show predictions based on market access and the residuals from equation (19). Red circles show predictions based on market access alone from equation (19). . For employment ($\log \widehat{L}_i$) and commercial floor space prices ($\log \widehat{q}_i$), model counterfactuals are exactly equal to the predictions based on market access and the residual (blue circles on the 45 degree line). For residents ($\log \widehat{R}_i$) and residential floor pries ($\log \widehat{Q}_i$), model counterfactual are approximately equal to the predictions based on market access and the residual (blue circles away from the 45 degree line), where the approximation is taken around an equilibrium with prohibitive travel costs.

On the horizontal axis of each panel, we show the true log changes in each variable from solving for the counterfactual equilibrium in the full non-linear model. On the vertical axis of each panel, we show the predicted log changes in each variable based on changes in market access ($\widehat{\Phi}_n^R$, $\widehat{\Phi}_n^L$) from the reduced-form system of equations (19). The blue circles (labelled total in the legend) show the overall predictions taking into account both changes in market access ($\widehat{\Phi}_n^R$, $\widehat{\Phi}_n^L$) and changes in the residuals ($\widehat{e}_n^L$, $\widehat{e}_n^q$, $\widehat{e}_n^R$, $\widehat{e}_n^Q$).[16] The red circles (labelled market access in

[16] With no-arbitrage between alternative uses of floor space, the true log changes in the price of commercial and

the legend) show the predictions based on changes in market access alone ($\widehat{\Phi}_n^R$, $\widehat{\Phi}_n^L$), where the changes in the residuals ($\widehat{e}_n^L$, $\widehat{e}_n^q$, $\widehat{e}_n^R$, $\widehat{e}_n^Q$) are set equal to zero.

For employment (top left) and commercial floor space prices (top right), the overall predictions based on changes in market access and changes in the residuals are necessarily equal to the true counterfactual changes, as reflected in the blue circles lieing along the 45-degree line. This pattern of results reflects the fact that the reduced-form relationships (19) hold exactly for employment and commercial floor space prices.

For residents (bottom left) and residential floor space prices (bottom right), these overall predictions can diverge from the true counterfactual changes, as reflected in the blue circles departing from the 45-degree line. This pattern of results reflects the fact that the reduced-form relationships (19) are only approximations around an equilibrium with prohibitive commuting costs for residents and residential floor space prices. Although the gap from the 45-degree line varies across locations, the magnitude of this variation is limited, except for a relatively small number of locations for residential floor space prices. In part, these results reflect the fact that we start from an initial equilibrium in which there is relatively little commuting, with more than 90 percent of workers living within 5km of their workplace. Therefore, the approximation around an initial equilibrium with prohibitive commuting cost is relatively good.

In contrast, we find that the predictions based on market access alone can diverge substantially from the true counterfactual changes, as reflected in the red circles departing substantially from the 45-degree line for all four variables. The extent of the error is not constant, but instead differs substantially across locations. The magnitude of the departure from the 45-degree line is greater for employment and commercial floor space prices than for residents and residential floor space prices, perhaps in part because commercial economic activity is more spatially concentrated across locations than residential economic activity.

To provide further evidence on the quantitative success of the predictions based on market access alone, we use the log linear structure of the reduced-form relationships (19). We regress the true counterfactual log changes in each variable on the log change predicted by either residence

residential floor space are equal to one another: $\log(\widehat{q}_n) = \log\left(\widehat{Q}_n\right)$. Nevertheless, the overall predictions for these variables based on changes in market access and the residuals need not equal one another, because the reduced-form system of equations (19) only holds as an approximation for residents and residential floor space prices.

Table 2: Market Access Predictions

	Regression Slope	R-squared
Employment ($\log\left(\widehat{L}_n\right)$)	1.216***	0.816
Commercial Floor Prices ($\log\left(\widehat{q}_n\right)$)	0.785***	0.649
Residents ($\log\left(\widehat{R}_n\right)$)	1.784***	0.721
Residential Floor Prices ($\log\left(\widehat{Q}_n\right)$)	0.608***	0.546

Note: Results of regressions of the log change in each variable between the counterfactual and initial equilibria from the solution of the non-linear model on the predicted change based on market access alone from equation (19) for our baseline specification without agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$). For employment ($\log \widehat{L}_i$) and commercial floor space prices ($\log \widehat{q}_i$), counterfactuals log changes in the model are exactly equal to the predictions based on market access and a residual. For residents ($\log \widehat{R}_i$) and residential floor pries ($\log \widehat{Q}_i$), counterfactual log changes in the model are approximately equal to the predictions based on market access and the residual, where the approximation is taken around an equilibrium with prohibitive travel costs.

($\log \widehat{\Phi}_n^R$) or workplace ($\log \widehat{\Phi}_n^L$) market access:

$$\begin{aligned}
\log \widehat{L}_n &= \varpi^L + \varsigma^L \left[\frac{1}{1+\epsilon\left(1-\beta\right)} \log \widehat{\Phi}_n^L\right] + u_n^L, \\
\log \widehat{q}_n &= \varpi^q + \varsigma^q \left[\frac{\beta}{1+\epsilon\left(1-\beta\right)} \log \widehat{\Phi}_n^L\right] + u_n^q, \\
\log \widehat{R}_n &= \varpi^R + \varsigma^R \left[\frac{\alpha}{1+\epsilon\left(1-\alpha\right)} \log \widehat{\Phi}_n^R\right] + u_n^R, \\
\log \widehat{Q}_n &= \varpi^Q + \varsigma^Q \left[\frac{1+\epsilon}{\epsilon\left(1+\epsilon\left(1-\alpha\right)\right)} \log \widehat{\Phi}_n^R\right] + u_n^Q,
\end{aligned} \tag{21}$$

where (ϖ^L, ϖ^q, ϖ^R, ϖ^Q) are the regression intercepts; (ς^L, ς^q, ς^R, ς^Q) are the regression slope coefficients; and (u_n^L, u_n^q, u_n^R, u_n^Q) are the regression residuals.

Table 2 reports the estimated slope coefficients and R-squared from these regressions. We find that market access is an imperfect predictor of the true counterfactual log change for all four variables. We obtain estimated slope coefficients that are substantially and statistically significantly different from one (ranging from 0.608 to 1.784). We find regression R-squared that can differ substantially from one (ranging from 0.546-0.816), highlighting that the residuals can account for a quantitatively relevant share of the overall variation in each variable.

Therefore, we find that predictions based on market access alone can be quite misleading for the reorganization of economic activity across locations in response to a transport improvement. Our baseline specification assumes no agglomeration forces, such that productivity and amenities are constant. As a result, the only source of error in predictions based on market access alone is the endogenous reallocation of floor space between commercial and residence use ($\widehat{H}_n^L \neq 1$ and $\widehat{H}_n^R \neq 1$). In Online Appendix D, we show that we find a similar pattern of results in our augmented specification incorporating agglomeration forces, in which the error term includes both changes in productivity and amenities ($\widehat{A}_n \neq 1$ and $\widehat{B}_n \neq 1$) and changes in the allocation of floor space between alternative uses ($\widehat{H}_n^L \neq 1$ and $\widehat{H}_n^R \neq 1$).

5.7 Aggregate Implications

We have so far highlighted that a key advantage of quantitative urban models is that they provide a framework for modelling the general equilibrium reorganization of economic activity in response to transport infrastructure improvements. This reorganization can be challenging to capture in conventional cost-benefit analyses, because they are typically partial equilibrium in nature. This reorganization also can be challenging to fully capture using conventional measures of market access, in the presence of endogenous reallocations of floor space between alternative uses and spillovers of agglomeration externalities across locations. While our quantitative results so far have been concerned with the spatial distribution of economic activity across locations within the city, we now compare alternative approaches for evaluating the aggregate impact of the transport improvement for the city as a whole.

Under our assumption of a closed city, the aggregate impact of the construction on the railway network on worker welfare is captured by the relative change in expected worker utility ($\widehat{U} = U'/U$). The extreme value distribution of idiosyncratic preferences implies that expected worker utility conditional on choosing a residence-workplace is the same across all residence-workplace pairs, and equal to expected utility for the city as a whole.

Under our assumption of local landlords who consume only the freely-traded numeraire good, the change in the welfare of local landlords is captured by the relative change in income from the ownership of local floor space: $\widehat{\mathbb{R}}_n = \mathbb{R}_n'/\mathbb{R}_n$,where $\mathbb{R}_n = Q_n H_n^R + q_n H_n^L$. Therefore, the construction of the railway network has distributional consequences across landlords, depending on the location in which they own floor space. We focus here on the railway network's impact on the aggregate income of all landlords, summing across all locations within the city: $\widehat{\mathbb{R}} = \mathbb{R}'/\mathbb{R}$,where $\mathbb{R} = \sum_{n \in \mathbb{N}} \left(Q_n H_n^R + q_n H_n^L \right)$).

We compare the aggregate impact of the construction of the railway network on worker and landlord welfare in our quantitative urban model to two alternative benchmarks. First, we examine the direct savings in travel costs, which are incurred in terms of utility. The savings in travel costs for each worker commuting from residence n to workplace i as a result of the transport improvement are $\widehat{\kappa}_{ni}^{-1} = \left(\kappa_{ni}' / \kappa_{ni} \right)^{-1}$. We compute the weighted average of these savings in travel costs across all bilateral residence-workplace pairs, using the commuting probabilities in the initial equilibrium as weights: $\sum_{n \in \mathbb{N}} \sum_{i \in \mathbb{N}} \lambda_{ni} \left(\widehat{\kappa}_{ni}^{-1} \right)$.

Second, a classic result from macroeconomics is Hulten's Theorem that for an efficient economy the aggregate welfare gain from a technological improvement can be captured up to a first-

Table 3: Aggregate Impacts of the Railway Network

	Without Agglomeration	With Agglomeration
Relative Worker Expected Utility	1.330	1.451
Relative Aggregate Value of Floor Space	1.118	1.176
Commuting Costs Saved (Initial Commuting Shares)	1.164	1.170
Commuting Costs Saved (Initial GDP Shares)	1.214	1.222

Note: first row shows relative increase in worker expected utility between the initial and counterfactual equilibria: $\widehat{U} = U'/U$; second row shows the relative increase in the aggregate value of floor space between the initial and counterfactual equilibria: $\left(\sum_{i \in \mathbb{N}} \left(Q_i' H_i^{R\prime} + q_i' H_i^{L\prime}\right)\right) / \left(\sum_{i \in \mathbb{N}} \left(Q_i H_i^R + q_i H_i^L\right)\right)$; third row is the weighted average saving in travel costs for each residence-workplace pair using initial commuting shares as weights: $\sum_{n \in \mathbb{N}} \sum_{i \in \mathbb{N}} \lambda_{ni} \widehat{\kappa}_{ni}^{-1}$; fourth row is the weighted average saving in travel costs for each residence-workplace pair using initial gross domestic product (GDP) shares as weights: $\sum_{n \in \mathbb{N}} \sum_{i \in \mathbb{N}} \left(w_i \lambda_{ni}/\beta\right) \widehat{\kappa}_{ni}^{-1}$; first column reports results for the specification without agglomeration forces ($\eta_A = \eta_B = \delta_A = \delta_B = 0$); second column reports results for the specification with agglomeration forces ($\eta_A = \eta_B = 0.10$, $\delta_A = \delta_B = 0.08$).

order approximation as the weighted average of the impact of the technological improvement at the disaggregated level. The relevant weights are the Domar weights in the initial equilibrium. In our closed-city with no intermediate inputs, these Domar weights equal shares of gross domestic product (GDP) in the initial equilibrium. GDP itself is defined as the sum of the income of workers and landlords across all locations. Using our assumption of a Cobb-Douglas production technology, we compute the GDP-share weighted average of the reduction in travel costs as follows: $\sum_{n \in \mathbb{N}} \sum_{i \in \mathbb{N}} \left(w_i \lambda_{ni}/\beta\right) \left(\widehat{\kappa}_{ni}^{-1}\right)$.

Table 3 summarizes these different measures of the aggregate impact of the construction of the railway network for our baseline specification without agglomeration forces (first column) and our augmented specification with agglomeration forces (second column). We find a counterfactual increase in worker expected utility of around 33 percent in our baseline specification (first row). This substantial welfare gain reflects the large counterfactual reduction in commuting costs. Travel times (τ_{ni}) fall by a multiple of 5 between grid points that are directly connected by a railway line, which leads to large reductions in commuting costs ($\kappa_{ni} = \exp\left(\kappa \tau_{ni}\right)$). The resulting gain in expected worker utility reflects the enhanced ability of workers to sort across residence-workplace pairs to take advantage of variation in production fundamentals, idiosyncratic preference draws, and bilateral travel costs.

We find a somewhat smaller, but still substantial, counterfactual increase in landlord income of 12 percent (second row). Landlords are geographically immobile and experience no idiosyncratic preference draws, which implies that the construction of the railway network does not directly affect their utility. Nevertheless, the enhanced ability of workers to sort across residence-workplace pairs to take advantage of differences in production fundamentals raises the aggregate value of land in our counterfactuals. Locations become increasingly able to specialize as workplaces and residences, to take advantage of the comparative advantage created by differences in production fundamentals and bilateral travel costs.

We find that the travel cost saved from the construction of the railway network provide an imperfect proxy for the increase in expected worker utility. Using initial commuting shares as weights, we find an average travel cost saved of 16 percent in our baseline specification (third row). This divergence between the gain in worker expected utility and average travel cost saved is perhaps unsurprising. In partial equilibrium cost-benefit analyses, travel cost saved only provides a first-order approximation to the true welfare effect. The construction of the railway network involves a large change in travel costs, with the commuting shares in the initial equilibrium differing substantially from those in the counterfactual equilibrium. Therefore, the second-order and higher terms in the Taylor-series expansion around the true welfare effect can be large.

Using initial GDP shares as weights, we find an average travel time saved of 21 percent in our baseline specification (fourth row), which is again substantially smaller than the increase in expected worker utility. Again this difference is perhaps unsurprising. Hulten's Theorem only holds exactly for an economy with a Cobb-Douglas network structure. Our model of worker commuting decisions features a constant elasticity of substitution (CES) gravity equation, with a commuting elasticity that differs from one. Therefore, Hulten's Theorem only provides a first-order approximation to the true welfare effects. The construction of the railway network involves a large change in travel costs, and the second-order and higher terms in the Taylor-series expansion around the true welfare effect can be large.

In our augmented specification incorporating agglomeration forces, we find larger increases in worker expected utility from the construction of the railway network (45 percent in the second column compared to 33 percent in the first column). As the reduction in commuting costs allows locations to specialize according to their comparative advantage, the concentration of employment in locations with high production fundamentals raises productivity in those locations through production externalities, thereby increasing welfare. Similarly, the concentration of residents in locations with comparative advantages as residences raises amenities through residential externalities, thereby increasing welfare.

Again we find discrepancies between the true increase in worker expected utility and the predictions from benchmarks based on travel cost saved (45 percent for worker expected utility compared to 17 and 22 percent for average travel cost saved using initial commuting shares and GDP shares as weights, respectively). Again these discrepancies reflect the fact that the construction of the railway network involves a large change in travel costs. Additionally, agglomeration forces in the form of production and residential externalities provide a source of market failure, such that Hulten's Theorem need no longer hold as a first-order approximation.

Overall, we find that the aggregate predictions of quantitative urban models can differ substantially from benchmarks based on travel cost saved for large changes in the transport network. Although the change in the transport network that we consider in our numerical example involves a substantial change in relative travel costs, it is not historically unprecedented. We calibrate the size of the reduction in travel costs to match empirical moments for changes in commuting distances following the construction of London's 19th-century railway network. There are several other historical examples of large-scale changes in transport technology, including for example the construction of the U.S. interstate highway network after the Second World War, and ongoing advances in the development of autonomous vehicles.

6 Conclusions

A major breakthrough in recent research is the development of quantitative urban models. These models are sufficiently rich to capture observed features of the data, such as many asymmetric locations and a rich geography of the transport network. Yet they remain sufficiently tractable as to permit an analytical characterization of their theoretical properties, such as the existence and uniqueness of the equilibrium. With only a small number of parameters to be estimated, these models lend themselves to transparent identification. Since they rationalize the observed distribution of economic activity in the data, they can be used to undertake counterfactuals for the impact of empirically-relevant public-policy interventions, such as the construction of railway line along a particular route.

We compare evaluations of a transport improvement using conventional cost-benefit analysis, sufficient statistics approaches based on changes in market access, and model-based counterfactuals. When the demand for travel is perfectly inelastic and the change in travel cost equals the saving in travel time multiplied by the value of time, the welfare gains from a transport improvement can be measured using conventional cost-benefit analysis as the total value of the travel time saved. More generally, this measure provides a first-order approximation to the welfare gains from a transport improvement, which can differ from its full general equilibrium impact for large changes in the transport network.

Within the class of quantitative urban models characterized by a gravity equation for commuting flows, the reorganization of economic activity in response to a transport improvement can be approximated using measures of residence and workplace market access. Residence market access captures proximity to surrounding sources of employment, while workplace market access captures proximity to surrounding sources of residents. This approximation is undertaken around an initial equilibrium with prohibitive commuting costs, and abstracts from changes in the supplies of residential and commercial floor space, and spillovers of production and residential externalities across locations.

In quantitative urban models, the spatial distribution of economic activity within cities is determined by the interaction between exogenous differences in location characteristics and endogenous agglomeration forces. Exogenous location characteristics include production fundamentals, residential fundamentals and position in geographical space. An improvement in the transport network reduces the costs to workers of separating their residence and workplace, which allows locations to increasingly specialize according to their comparative advantage in residential and commercial activity. In the presence of agglomeration forces, the increased concentration of employment and residents across locations magnifies exogenous differences in production and residential fundamentals, and leads to further increases in the specialization of locations as residences and workplaces.

We illustrate the use of quantitative urban models to evaluate a transport improvement using a numerical example of a city. By focusing on a numerical example, we consider a setting in which we know the true data generating process (DGP) and model parameters. Therefore, the data are generated according to the model, and we can examine the success of alternative approaches to approximating the true impact of the transport improvement. We show that the direct impact of a transport improvement on travel time can be quite misleading for the general equilibrium impact, once the reorganization of economic activity in response to the new transport technology is taken into account. We show that the predicted reorganization of economic activity based on changes in market access can differ substantially from the predicted reorganization from model-based counterfactuals for large changes in the transport network.

References

Ahlfeldt, G. and E. Pietrostefani (2019): "The Economic Effects of Density: A Synthesis," *Journal of Urban Economics*, 111, 93–107.

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

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

——— (2022a): "Supply and Demand in Space," *Journal of Economic Perspectives*, 37, 3–28.

——— (2022b): "The Welfare Effects of Transportation Infrastructure Improvements," *Review of Economic Studies*, 89, 2911–2957.

——— (2025): "Quantitative Economic Geography," in *Handbook of Regional and Urban Economics*, ed. by D. Donaldson and S. J. Redding, Amsterdam: Elsevier North Holland, vol. 6.

Allen, T., C. Arkolakis, and X. Li (2024): "On the Equilibrium Properties of Spatial Models," *American Economic Review: Insights*, 6, 472–89.

Almagro, M., F. Barbieri, J. C. Castillo, N. Hickok, and T. Salz (2024): "Optimal Urban Transportation Policy: Evidence from Chicago," *NBER Working Paper*, 32185.

Almagro, M. and T. Domínguez-lino (2024): "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," *NBER Working Paper*, 32304.

Alonso, W. (1964): *Location and Land Use*, Cambridge MA: Harvard.

Baum-Snow, N. (2007): "Did Highways Cause Suburbanization?" *Quarterly Journal of Economics*, 122, 775–780.

Baum-Snow, N. (2020): "Urban Transport Expansions and Changes in the Spatial Structure of US Cities: Implications for Productivity and Welfare," *Review of Economics and Statistics*, 102, 929–945.

Baum-Snow, N., L. Brandt, J. Henderson, M. A. Turner, and Q. Zhang (2017): "Roads, Railroads and Decentralization of Chinese Cities," *Review of Economics and Statistics*, 99, 435–448.

Baum-Snow, N., V. Henderson, M. A. Turner, Q. Zhang, and L. Brandt (2020): "Does Investment in National Highways Help or Hurt Hinterland City Growth?" *Journal of Urban Economics*, 115, 103124.

Baum-Snow, N. and M. E. Kahn (2005): "Effects of Urban Rail Transit Expansions: Evidence from Sixteen Cities, 1970-2000," *Brookings-Wharton Papers on Urban Affaires: 2005*, 1, 147–197.

Baum Snow, N. and H. Lu (2024): "The Microgeography of Housing Supply," *Journal of Political Economy*, 132, 1793–2178.

Billings, S. B. (2011): "Estimating the Value of a New Transit Option," *Regional Science and Urban Economics*, 41, 525–536.

Bordeu, O. (2023): "Commuting Infrastructure in Fragmented Cities," Chicago Booth School of Business.

Brooks, L. and B. Lutz (2018): "Vestiges of Transit: Urban Persistence at a Micro Scale," *Review of Economics and Statistics*, 101, 385–399.

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2018): "The Impact of Regional and Sectoral Productivity Changes on the U.S. Economy," *Review of Economic Studies*, 85, 2042–2096.

Chandra, A. and E. Thompson (2000): "Does Public Infrastructure affect Economic Activity? Evidence from the Rural Interstate Highway System," *Regional Science and Urban Economics*, 30, 457–490.

Chetty, R. (2009): "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1, 451–87.

Costinot, A. and A. Rodríguez-Clare (2014): "Trade Theory with Numbers: Quantifying the Consequences of Globalization," in *Handbook of International Economics*, ed. by E. Helpman, G. Gopinath, and K. Rogoff, Amsterdam: North-Holland, vol. 4, chap. 4, 197–261.

Davis, M. A. and F. Ortalo-Magné (2011): "Housing Expenditures, Wages, Rents," *Review of Economic Dynamics*, 14, 248–261.

Dekle, R., J. Eaton, and S. Kortum (2007): "Unbalanced Trade," *American Economic Review*, 97, 351–355.

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

Dingel, J. I. and F. Tintelnot (2020): "Spatial Economics for Granular Settings," *NBER Working Paper*, 27287.

Donaldson, D. (2018): "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, 108, 899–934.

——— (2025): “Transportation,” in *Handbook of Regional and Urban Economics*, ed. by D. Donaldson and S. J. Redding, Amsterdam: Elsevier North Holland, vol. 4, chap. 50.

Donaldson, D. and R. Hornbeck (2016): “Railroads and American Economic Growth: A Market Access Approach,” *Quarterly Journal of Economics*, 131, 799–858.

Duranton, G., P. Morrow, and M. A. Turner (2014): “Roads and Trade: Evidence from the US,” *Review of Economic Studies*, 81, 681–724.

Duranton, G. and M. A. Turner (2011): “The Fundamental Law of Road Congestion: Evidence from US Cities,” *American Economic Review*, 101, 2616–52.

——— (2012): “Urban Growth and Transportation,” *Review of Economic Studies*, 79, 1407–1440.

Faber, B. (2014): “Integration and the Periphery: The Unintended Effects of New Highways in a Developing Country,” *Review of Economic Studies*, 81, 1046–1070.

Fajgelbaum, P. and C. Gaubert (2020): “Optimal Spatial Policies, Geography, and Sorting,” *Quarterly Journal of Economics*, 135, 959–1036.

Fajgelbaum, P. D. and E. Schaal (2020): “Optimal Transport Networks in Spatial Equilibrium,” *Econometrica*, 88, 1411–1452.

Fogel, R. W. (1964): *Railroads and American Economic Growth: Essays in Econometric History*, Baltimore, MD: Johns Hopkins University Press.

Fortheringham, S. and M. O’Kelly (1989): *Spatial Interaction Models: Formulations and Applications*, Dordrecht: Kluwer.

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

Fujita, M. and H. Ogawa (1982): “Multiple Equilibria and Structural Transformation of Non-monocentric Urban Configurations,” *Regional Science and Urban Economics*, 12, 161–196.

Fujita, M. and J.-F. Thisse (2002): *Economics of Agglomeration: Cities, Industrial Location and Regional Growth*, Cambridge: Cambridge University Press.

Fukui, M., E. Donald, and Y. Miyauchi (2024): “Unpacking Welfare in a Spatial Economy,” Boston University, mimeograph.

Gechter, M. and N. Tsivanidis (2022): “Spatial Spillovers from Urban Redevelopments: Evidence from Mumbai’s Textile Mills,” University of California, Berkeley.

Gibbons, S. and S. Machin (2005): "Valuing Rail Access Using Transport Innovations," *Journal of Urban Economics*, 57, 148–169.

Gonzalez-Navarro, M. and M. A. Turner (2018): "Subways and Urban Growth: Evidence from Earth," *Journal of Urban Economics*, 108, 85–106.

Greenstone, M., R. Hornbeck, and E. Moretti (2010): "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings," *Journal of Political Economy*, 118, 536–598.

Heblich, S., S. Redding, and D. Sturm (2020): "The Making of the Modern Metropolis: Evidence from London," *Quarterly Journal of Economics*, 135, 2059–2133.

Hornbeck, R. and M. Rotemberg (2024): "Growth Off the Rails: Aggregate Productivity Growth in Distorted Economies," *Journal of Political Economy*, 132, 3547–3602.

Jackson, K. T. (1987): *Crabgrass Frontier: The Suburbanization of the United States*, Oxford: Oxford University Press.

Jiang, Y. (2025): "Optimal Policy with Strategic Transportation Investment," Princeton University.

Jones, I. S. (1977): *Urban Transport Appraisal*, London: John Wiley and Sons.

Kline, P. and E. Moretti (2014): "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *Quarterly Journal of Economics*, 129, 275–331.

Kreindler, G. E., A. Gaduh, T. Graff, R. Hanna, and B. A. Olken (2023): "Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta," *NBER Working Paper*, 31369.

Kreindler, G. E. and Y. Miyauchi (2023): "Measuring Commuting and Economic Activity inside Cities with Cell Phone Records," *Review of Economics and Statistics*, 105, 899–909.

Lucas, R. E. and E. Rossi-Hansberg (2002): "On the Internal Structure of Cities," *Econometrica*, 70, 1445–1476.

McDonald, J. and D. McMillen (2010): *Urban Economics and Real Estate: Theory and Policy*, Hoboken, NJ: John Wiley & Sons.

McDonald, J. F. and C. I. Osuji (1995): "The Effect of Anticipated Transportation Improvement on Residential Land Values," *Regional Science and Urban Economics*, 25, 261–278.

McFadden, D. (1974): "The Measurement of Urban Travel Demand," *Journal of Public Economics*, 3, 303–328.

Melo, P. C., D. J. Graham, and R. B. Noland (2009): "A Meta-analysis of Estimates of Urban Agglomeration Economies," *Regional Science and Urban Economics*, 39, 332–342.

Michaels, G. (2008): "The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System," *Review of Economics and Statistics*, 90, 683–701.

Mills, E. S. (1967): "An Aggregative Model of Resource Allocation in a Metropolitan Centre," *American Economic Review*, 57, 197–210.

Miyauchi, Y., K. Nakajima, and S. Redding (2022): "The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data," *NBER Working Paper*, 28497.

Mongey, S. and M. Waugh (2024): "Discrete Choice, Complete Markets and Equilibrium," *NBER Working Paper*, 32135.

Monte, F., C. Porcher, and E. Rossi-Hansberg (2023): "Remote Work and City Structure," *NBER Working Paper*, 31494.

Monte, F., S. Redding, and E. Rossi-Hansberg (2018): "Commuting, Migration and Local Employment Elasticities," *American Economic Review*, 108, 3855–3890.

Muth, R. (1969): *Cities and Housing*, Chicago: University of Chicago Press.

Nagy, D. K. (2023): "Hinterlands, City formation and Growth: Evidence from the U.S. Westward Expansion," *Review of Economic Studies*, 90, 3238–3281.

Owens III, R., E. Rossi-Hansberg, and P.-D. Sarte (2020): "Rethinking Detroit," *American Economic Journal*, 12, 258–305.

Redding, S. J. (2016): "Goods Trade, Factor Mobility and Welfare," *Journal of International Economics*, 101, 148–167.

——— (2023): "Quantitative Urban Models: From Theory to Data," *Journal of Economic Perspectives*, 37, 75–98.

——— (2025): "Quantitative Urban Economics," in *Handbook of Regional and Urban Economics*, ed. by D. Donaldson and S. J. Redding, Amsterdam: Elsevier.

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

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

——— (2024): "Neighborhood Effects: Evidence from Wartime Destruction in London," *NBER Working Paper*, 32333.

Redding, S. J. and M. A. Turner (2015): "Transportation Costs and the Spatial Organization of Economic Activity," in *Handbook of Regional and Urban Economics*, ed. by G. Duration, J. V. Henderson, and W. Strange, Amsterdam: Elsevier, chap. 20, 1339–1398.

Rosenthal, S. S. and W. C. Strange (2004): "Evidence on the Nature and Sources of Agglomeration Economics," in *Handbook of Regional and Urban Economics*, ed. by J. V. Henderson and J. Thisse, Amsterdam: Elsevier, vol. 4.

Saiz, A. (2010): "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 125, 1253–1296.

Severen, C. (2023): "Commuting, Labor and Housing Effects of Mass Transportation: Welfare and Identification," *Review of Economics and Statistics*, 105, 1073–1091.

Tsivanidis, N. (2024): "The Aggregate And Distributional Effects Of Urban Transit Infrastructure: Evidence From Bogotá's TransMilenio," *American Economic Review*, forthcoming.

Valentinyi, A. and B. Herrendorf (2008): "Measuring Factor Income Shares at the Sectoral Level," *Review of Economic Dynamics*, 11, 820–835.

Warner, S. B. (1978): *Streetcar Suburbs: The Process of Growth in Boston, 1870-1900*, Cambridge MA: Harvard University Press.

Weiwu, L. (2024): "Unequal Access: Racial Segregation and the Distributional Impacts of Interstate Highways in Cities," MIT, mimeograph.