The Impact of Tokyo's Train System*

Mark Bamba

Dec 30, 2024

Abstract

We evaluate the impact of Tokyo's train system on the spatial distribution of economic activity and welfare. Since its first railway was built in 1872, Tokyo's rail system has grown to become one of the world's most heavily utilized urban railway network, serving about half of of Tokyo's roughly 40 million residents daily. Using an urban model of commuting that incorporates workers' choice of residence location, work location, and mode of transport (between train, car, walking, cycling, and bus), we analyze the impact of Tokyo's train system by simulating the spatial distribution of Tokyo's economy if it no longer had a train system and automobiles were instead the primary mode of transport (which accounted for $\sim$ 23% of commuting trips in 2018). Simulations that keep today's road infrastructure, road congestion, floorspace, and total population fixed imply that central Tokyo, which relies heavily on the train system, would lose 3% of its population and 12 % in employment while rural areas that are already car-dependent gain 8% in population and 20% in employment. Welfare declines by 33%. Preliminary simulations that incorporate the impact of road congestion on car travel times imply that population and employment decline by 9% and 15% respectively in the city center while the peripheries gain 13% in population and 32% in employment. While increased congestion accounts for a 12% increase in automobile commute times on average (from 22 to 26 min), the spatial reorganization to shorter commutes implies that the average automobile commute time decreases by 2 minutes, relative to the current 2020 equilibrium. Welfare declines by an additional 4%. Our findings show quantitatively how automobile-focused urban areas lead to low-density "sprawl" as is observed in many American cities.

*This project would not have been possible without Eduardo Morales and Steve Redding's continued advice and support. Atsushi Yamagishi and Kentaro Nakajima provided invaluable guidance throughout. This project has benefitted from conversations with David Weinstein, Richard Rogerson, Owen Zidar, Derek Wenning, Michael Jenuwine, Pietro Buri, Mariko Nakagawa, Shota Fujishima, Chihiro Shimizu, Yuyang Jiang among others. Participants from the Princeton Student Trade Workshop and Hitotsubashi's Urban Seminar are thanked for their constructive comments. This project is the result of a collaboration with the Center for Spatial Information Science, the University of Tokyo (Project No. 1296) and used the 1991 and 1996 *Establishment and Enterprise Census Statistical Data by Small Areas*, and the 1990 and 1995 *Population Census Tabulation for Small Areas*, all of which were provided by Sinfonica.

1 Introduction

What are the economic gains of urban public transit systems? As the urban population is estimated to grow worldwide from the current 56% of the world's population to 70% by 2050, there is a question of what the economic benefits to transportation infrastructure projects are in order for governments to determine which ones to implement and how.

This paper studies the economic impact of the public transit system in Tokyo, Japan, one of the largest metropolitan areas in the world and host to one of the largest urban railway networks in the world. Since the first railway in Tokyo was built in 1872, the network has grown to some 185 lines facilitating approximately 40 million passenger trips daily within the metropolitan area. Tokyo meanwhile has grown to become the largest urban economy in the world by population, with 39 million residents in the metro area, and one of the largest by GDP, which was 1.8 trillion US dollars in 2019, or roughly 40 percent of Japan's GDP.[1]

In order to evaluate the impact of the entire train system, we employ an urban model of commuting that incorporates workers' choice of residence location, work location, and mode of transport (between train, car, walking, cycling, and bus) and simulate the spatial distribution of Tokyo's economy for a counterfactual where the city no longer had a train system, and automobiles were instead the primary mode of transport (which accounted for $sim23\%$ of commuting trips in 2018). The model is estimated and calibrated using population and employment data from Japan's censuses, land-price data from Japan's Official Land-Price Announcements, and commuting-flow data from Tokyo's Person-Trip Survey, together with parameters from existing literature.

Simulations that keep today's road infrastructure, road congestion, floorspace, and total population fixed predict that Tokyo's suburbs that rely heavily on the train system would lose 3% of its population and 12 % in employment while rural areas that are already car-dependent gain 8% in population and 20% in employment. Welfare declines by 33%. Preliminary simulations that incorporate the impact of road congestion on car travel times imply that population and employment decline by 9% and 15% respectively in the city center while the peripheries gain 13% in population and 32% in employment. While increased congestion accounts for an additional 12% in automobile commute time on average (from 22 to 26 min), the spatial reorganization of the economy implies that the average automobile commute time decreases by 2 minutes, relative to the current 2020 equilibrium. Welfare declines by an additional 4%. Our findings demonstrate quantitatively how automobile-focused urban areas lead to low-density "sprawl" as is observed in many American cities.

The following sections describe the historical background of Tokyo's train system, the data sources, the urban model of commuting, an extension that incorporates automobile road congestion, our estimation strategy, and results.

The appendix describes our estimation strategy for estimating spillover parameters, which are an important component of urban models to rationalize the observed density of economic activity in cities, as emphasized by the previous literature. Using the expansion of the Tokyo's train system

[1] In 2015 US dollars, PPP adjusted, according to the OECD's functional urban areas definition

from 1995-2020 as a natural experiment, we estimate how population and employment changes induced by the train expansion affected productivity and amenities across Tokyo. In future versions of this paper, we plan to incorporate these estimates into the model. For now, we present the results of the counterfactual simulations using parameters from the literature.

2 Related Literature

This paper contributes to several literatures.

First, the paper draws extensively on the recent literature on quantitative spatial economic models and their applications. These papers include Ahlfeldt et al. 2015, Tsivanidis 2023, Fajgelbaum and Schaal 2020, Allen and Arkolakis 2022, Monte, S. J. Redding, and Rossi-Hansberg 2018, S. J. Redding and Sturm 2008, Donaldson 2018, Donaldson and Hornbeck 2016, and Fajgelbaum, Gaubert, et al. 2023.

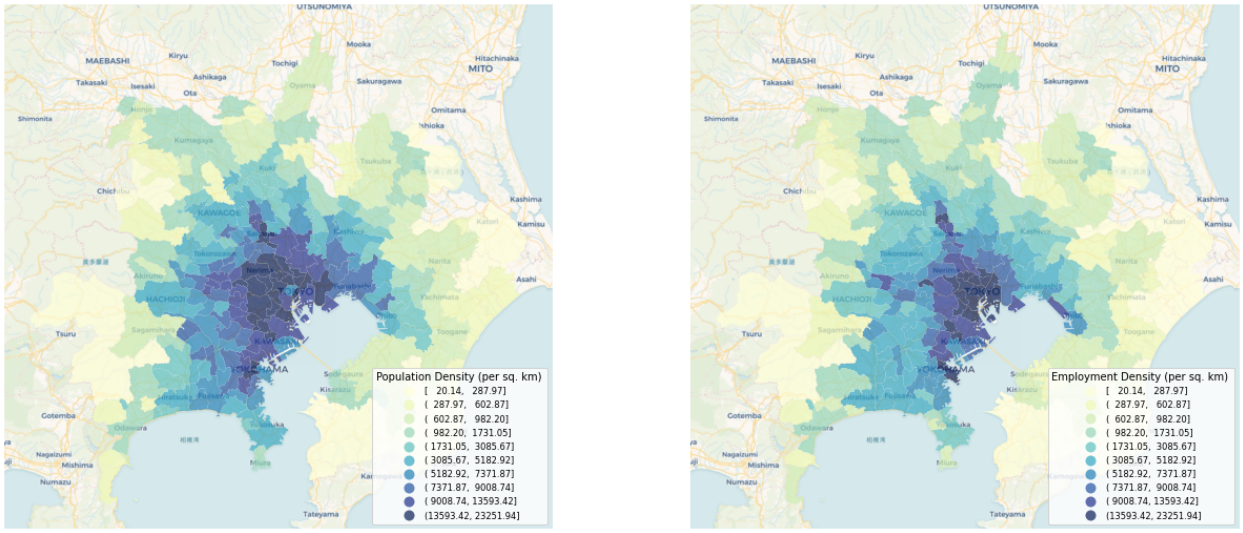
This paper applies the methods introduced in Ahlfeldt et al. 2015, Tsivanidis 2023, and Allen and Arkolakis 2022 to study the impact of an entire train system in the presence of automobile congestion.

Second, this paper relates to the literature on the economic impact of transportation infrastructure. These papers include Duranton, Nagpal, and M. Turner 2020, Baum-Snow 2007, Baum-Snow, Brandt, et al. 2017, Asahi, Herrera, and Silva 2023, Berger and Enflo 2017, Banerjee, Duflo, and Qian 2020, Faber 2014, and Brooks and Liscow 2021. Studies on Japan include Bernard, Moxnes, and Saito 2019, which studies how high-speed passenger rail improved firms' buyer-seller linkages, and Kazunobu et al. 2021, which studies the impact of high-speed rail on the distribution of economic activity across Japan. Relative to these studies, this paper provides the first model-consistent estimates of the economic impact of the train network in the world's largest urban economy. In doing so, we provide insights not only for Tokyo, but also for other cities around the world that are expanding their public transit systems as their urban populations grow.

Last, this paper specifically speaks to the literature on the economic impact of urban rail. Beginning with the classic text of John R. Meyer et al. 1959 and John Robert Meyer 1965, a large literature has documented that since the advent of the automobile and its subsequent ubiquity led to a decline in population density in urban areas, urban train systems have become economically unviable given their large fixed costs. More recent works by Baum-Snow, Kahn, and Voith 2005, Gonzalez-Navarro and M. A. Turner 2018, E. L. Glaeser and Kahn 2001, E. L. Glaeser and Kahn 2004, and E. Glaeser, Kahn, and Rappaport 2008 have also argued in a similar vein. This paper contributes to this literature by analyzing an important exception to the rule that urban rail systems are economically unviable, and seeking to understand the reason for this exception in the largest urban economy in the world that happens to have one of the largest urban rail systems in the world.

3 Historical Background

As of 2023, the prefecture of Tokyo, an administrative area of 23 wards plus some 39 neighboring townships and islands, had a population of 14.1 million people, roughly 11 percent of Japan's population, making it one of the most populous cities in the world. However, if it is defined based on its urban region, the city of Tokyo, also known as the *Greater Tokyo Area* or simply the *Tokyo metro*, extends to the the prefecture's neighoring prefectures such as Kanagawa, Saitama, and Chiba and has a population of 38 million people as of 2020, making it the most populous metropolitan area in the world.[2] Figure 1 shows the residential population density and employment density across the metro, a term we will henceforth use interchangeably with the term "Tokyo." Central Tokyo, shown in the map as the two densist regions, is roughly the region of the old capital, Edo, which expanded as the city's population grew between the 17th and 20th centuries.



Population Density Employment Density

Figure 1: Population and Employment per square kilometer in the Tokyo metro.

Edo was originally a fishing village that rose to prominence in 1603 when it was made the capital of the Tokugawa shogunate, the military (i.e. samurai) government of Japan that was in power between 1603 and 1868. The city grew at a steady pace of of about 1.7 percent per year (compound annual growht rate), from a population of about 10,000 in 1550 to roughly 1.4 million by 1850. Population growth accelerated following the Meiji Restoration in 1868, when the military government lost its power and the emperor's political supremacy was reinstated. As Japan opened its borders to trade with the west and the economy grew rapidly, the population grew by about 3

[2]The term "Daitoshiken," meaning the "metropolitan area of a metropolis" was first established in the 1960 national census, and the definition since the 1975 national census includes the city center, consisting of Tokyo's special wards and designated cities, as well as cities and towns that are connected to the central city with a commuting and school-going population aged 15 and above making up 1.5 percent or more of the resident population, along with cities and towns surrounding these areas, forming the "surrounding municipalities." We will use this definition of Tokyo for the purpose of this paper.

percent per year to roughly 5 million by 1890 and then 2 percent per year to 13 million by 1941 when Japan entered World War II. [3]

As part of its modernization effort, Japan built its first railway line in 1872, some 40 years after the first fully steam-powered public railway connected Liverpool and Manchester in Great Britain. Using British financing and over a hundred British and European technical consultants, the government of Japan undertoook construction for a train line to connect Shimbashi, in central Tokyo, and Yokohama, a city 27 km to the south that is now part of the Greater Tokyo Area (See Figure 2) (Free 2017, 64).

Soon, the Meiji government and the private sector began a public-private cooperation leading to rapid growth in the rail network across Japan, and by the year 1900, it spanned 4,674 kilometers across Japan (Hirooka 2000).

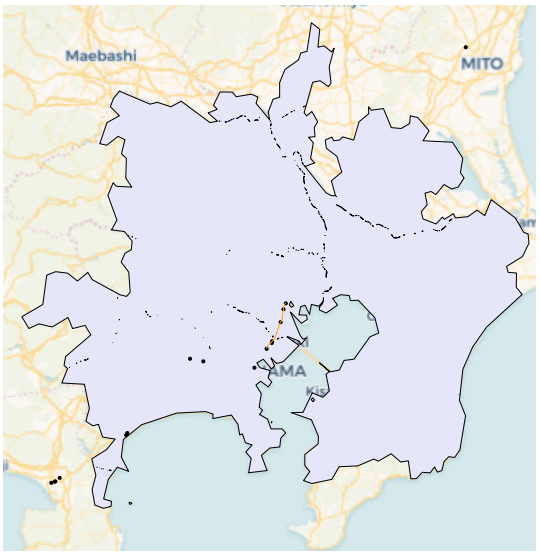
Within central Tokyo, the transit system between the 1910s to 1950s came to be dominated by trams, which operated on the same roads as automobiles. Tokyo's first subway line (i.e. underground train) was built in 1927, but due to the city's reliance on trams, the subway system did not grow rapidly until the post-war period, during which increased road congestion from the popularity of cars led to a rapid shift in usage from trams to subways. By the late 1960s, the city's tram system all but disappeared while the subway system had expanded to seven lines.

Figure 2 shows the evolution of Tokyo's train network from it's early years in 1880 to the present day. The train network has grown from a few lines in the city center to a comprehensive network that spans the entire metropolitan area. In 2020, there were 182 train lines in the Greater Tokyo Area spanning approximately 5200 km, more than the distance from Los Angeles to New York. These lines are operated by about 48 rail-trains operators, some of which are or were once owned by the government, but the majority of which are private companies.

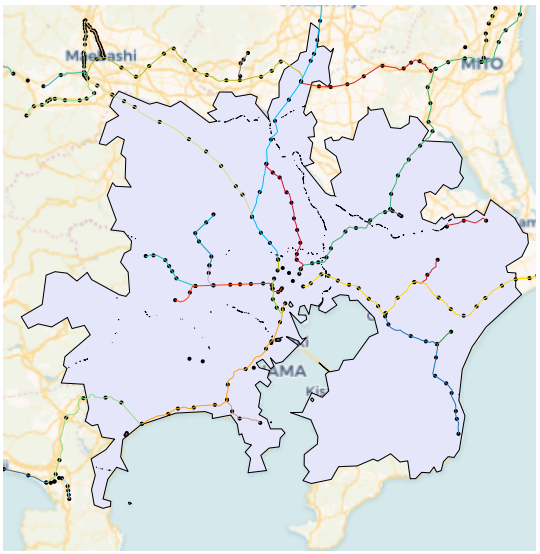
Over it's 150-year history, much of the construction costs have been paid for by national and local governments (through e.g. zero-interest loans, direct subsidies, subsidized bonds, and public investment) as well as through operators' retained earnings from fare revenue. However, a substantial portion has also been paid through so-called *land-value capture* schemes, where the increase in land values around train stations pay for the construction costs.[4] Each rail project has a different scheme with different stakeholders (public vs private rail operators, local vs national governments, different kinds of realestate developers), different methods of extracting land-value (e.g. land sales, land leases, development rights, etc), and different geography (e.g. urban vs rural, high vs low density, whether the rights of way can easily be acquired), schemes in Japan are typically characterized by substantial urban development around train stations aided by changes to minimum floor-area-ratios and other regulations to allow for denser development. Japan's rail companies are also known to operate large businesses that complement their rail business, such as in real-estate, shopping malls, restaurants, and hotels that are built surrounding train stations (for example,

[3]See for example Bernhofen and Brown 2005 on the impact of Japan's opening up of trade following the Meiji Restoration.

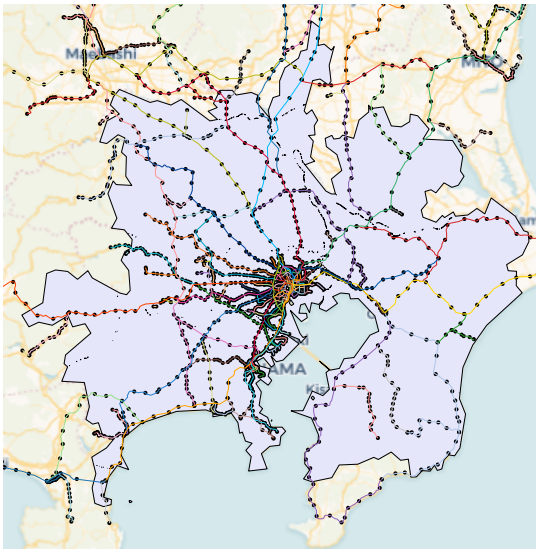
[4]If governmentm subsidies are paid for by taxes that originate from increased land-values (e.g. property and land taxes), then these schemes can also be viewed as a form of land-value capture. However, here we refer to other methods used to extract land-value for the purpose of paying for or reducing construction costs



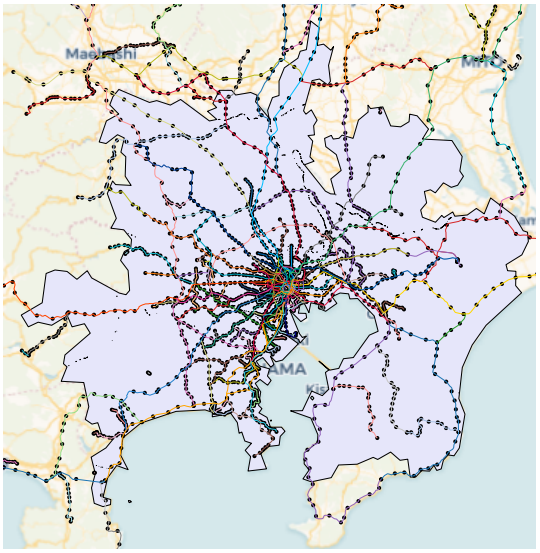
1880



1900



1950



2020

Figure 2: Train Expansion from 1880-2020

Colored lines are train lines while black dots are stations. Shaded region is the Greater Tokyo Area. Data is from the Ministry of Land, Infrastructure, Transport, and Tourism (MLIT)

"Real-estate and Hotels" and "Retail and Services" segments of JR East, Japan's largest train operator, made up 45 percent of its profits in FY2024). Since Henry George's observation that the benefits of public-good provisions are capitalized into land-rents, land-value capture schemes have been seen as an effective way to fund rail projects. In future work, we seek to analyze these schemes for the success of Tokyo's train network. In the current draft, we abstract from the sources of funding and instead focus on the impact of the train network on the distribution of economic activity absent the costs of construction. In future work, we intend to compare the benefits of the train system to to estimates of the costs of construction and maintenence of the system.

4 Data

We use a unit of analysis known as *oaza* in Japanese that is roughly the same size as a US census tract (and hence we will commonly use the term "census tract"), with an average population of approximately 2000 residents. We use this level of aggregation because most counterfactuals of interest (e.g. removing the train system) would have little to no change to commuting costs for commutes within a census tract.[5] If our analysis were done at a higher level of aggregation (e.g. municipality level), there would be a need to incorporate changes to commuting costs within a region. Another reason for this unit of analysis is that it is likely large enough that the model's assumption that agglomeration externalities do not spill over to neighboring districts is reasonable.[6] In addition, with approximately 8000 census tracts in our dataset and thus ∼64 million bilateral pairs, the computational burden should be manageable even with conventional computers. This draft presents its baseline counterfactual results at the census-tract level while the results with congestion are presented at the municipality level. Future versions of this paper will present results at the census-tract level for all counterfactuals.

There are four main datasets that are used to calibrate the quantitative model, in addition to several supplementary data sources that are used for estimation, calibration, and validation of the model. We describe each data set in turn and include additional details in the appendix.

First, population data by census tract comes from the Japanese Population Census. The census has been conducted almost every five years since 1920 in order to provide statistics on the state of Japan's population and households. This data is available at the block level from 1995, which is aggregated to the census tract level for our analysis.

Second, employment by region comes from two sources – the Economic Census for Business Activity, which has surveyed all establishments and companies every five years since 2012, and the Establishment and Enterprise Census, which occurred every five years before it was replaced by

[5]The average time (across origin census tracts) to travel to the closest census tract is 19 minutes, 5 minutes, 11 minutes, 5 minutes, and 6 minutes for train, car, bus, bicycle, and walking, respectively.

[6]Ahlfeldt et al. 2015 who use block-level data to estimate the strength of spillovers to neighboring blocks find rapid decay in spillovers that amount to a half-life of roughly two minutes for production externalities and one minute for residential externalities. The average travel time between neighboring census tracts in Tokyo is approximately 6 minutes by foot and therefore agglomeration spillover to neighboring tracts are likely small.

the Economic Census for Business Activity. These datasets, which are available at the block level from 1996, are also aggregated to the census-tract level for our analysis.[7]

For the purpose of our analysis, we define our "post"-expansion period to include the 2020 population and 2021 employment data, and define our "pre"-expansion period to include the 1995 population and 1996 employment data. The post-expansion datasets were obtained from Japan's portal site for official statistics, e-Stat, while data for the pre-expansion period were obtained from the University of Tokyo's Center for Spatial Information Science, which provides researchers with digitized versions of various censuses.

Data from the post-expansion period are used to calibrate the model, while data from the pre-expansion provide time variation to estimate the spillover parameters, as is discussed in section A.

Third, the primary source of land price data is the Official Land Price Announcements, which is published annually by the Japanese Ministry of Land, Infrastructure, Transport and Tourism (MLIT). Land prices across Japan are collected by real estate appraisers hired by the national government to provide policymakers with reliable land values that are important for calibrating national and local property tax levies. This data is supplemented by the prefecture-level surveys that are conducted similarly by local governments. Merging both datasets results in a panel of 11288 land plots across the Greater Tokyo Area in 2020. The full dataset is from the years 1983-2020 while we utilize the time variation between 1995 and 2020.

The national and prefectural governments select land plots for appraisal that are representative of a land-use category (residential, commercial, or industrial) at the municipality level (which has a population of approximately 100 thousand on average), and are typically tracked over time, unless the government deems the plot to no longer be representative of its category in a municipality.[8] Official municipality-level land-price indices then take a simple average of the land plots across a municipality, a methodology that we adopt to calculate the average municipality-level land price (used in our estimates of the model with congestion). For our census-tract-level analysis, our baseline estimates will also use the simple average of the land plots within a census tract.[9]

To convert land prices to floor prices, we follow Ahlfeldt et al. 2015 and use the assumption of a competitive construction sector with a Cobb–Douglas technology, as discussed in Section 5. Additional information on the land price data can be found in the appendix.

First, we calculate travel times for cars, cycling, and walking between origin-destination pairs using OpenRoute Services, an opensource tool that utilizes the street network provided by OpenStreetMap to calculate the shortest routes between locations and their associated travel times. The travel times are calculated between the centroids of census tracts. Since the exact centroid is typically not directly on the street network, OpenRoute Services determines the point on the street network that is closest to the centroid and uses this point as the origin or destination.

[7]The 1995 Establishment and Enterprise Census block-level data was digitized and provided to us through the University of Tokyo's Center for Spatial Information Science.

[8]See LaPoint 2020 for additional details on the land appraisal method.

[9]Variation in log land prices across municipalities account for 84 percent of the total variance in log land prices across plots. At the *oaza* level, the figure is 93 percent.

Next, we utilize data on the train network provided by MLIT. The data includes geographical information systems (GIS) shapefiles for the train lines and train stations, in addition to the beginning and ending dates of operation of segments of train lines from the 1800s to the present. We exclude high-speed rail lines (shinkansen) from our analysis, as they account for a share of commuting.[10]

We use the GIS shapefiles to construct a network on NetworkX, a network analysis library in Python and construct minimum travel times between the centroids of origin-destination pairs for trains. Train travel times between two train stations are calculated by dividing the distance of track between them by an assumed speed. Track lengths are inferred from the GIS data while determining train speed requires additional information. The average train speed for 24 of the major train lines were found online while the remaining train lines were assumed to travel at 40 km/hr, which was the approximate median travel speed of the train lines that we found. Passengers can transfer between train lines, but there is a 4-minute penalty for each transfer in addition to the time it takes to walk, which is assumed to occur at 4 km/hr. Travel times from the centroid of the census tract to the nearest train stations are computed using ORS assuming that they travel at half the speed of a car.

Additional details on these calculations can be found in section B.2.[11]

We use Tokyo's Personal Trip Survey for data on commuter flows and travel times between origin and destination districts by mode. The survey is of roughly 300,000 households (about 700,000 individuals) in the Greater Tokyo Area. Because commuting districts in this survey are larger than the census tracts, additional assumptions are needed to estimate the commuting elasticity, as will be described in section 7.2.

The data provides commuter flows and travel times by mode, but commuters often take multiple modes to reach their destination. The survey therefore uses a hierarchy of modes to determine the "primary mode of transport" for each trip. The hierarchy is as follows: train, bus, car, bicycle, and then foot. Thus, if a trip involves taking a train and a bus, the trip is classified as a train trip for the *primary mode*. [12] Based on these definitions, the survey finds that 54 percent of trips are by train, 23 percent by car, 10 percent by bicycle, 7 percent by foot, 2 percent by bus, and 3 percent by other modes. Other information on commuting in Tokyo can be found in Table 1. We use these definitions of modes as given from the travel survey in our gravity estimation and model calibration.

Travel times reported in the survey are those reported by commuters as the actual travel times between their home and workplace. We believe these times are more accurate than those that were computed with the GIS tools and subject to less systematic bias. We therefore calibrate our

[10] See Bernard, Moxnes, and Saito 2019 for a study on the impact of high-speed rail on firms' buyer-seller linkages and Kazunobu et al. 2021 for a study on the impact of high-speed rail on the distribution of economic activity across Japan.

[11] Services that calculate travel times between bilateral pairs of locations, such as Google Routes API, may be useful for calculating the travel times between locations in the current equilibrium. However, these tools are not readily available for public-transport routes in Tokyo and therefore are not used.

[12] For trips involving train, the survey found that travel between the origin or destination to the train station occurs 80 percent by foot, 8 percent by bus, 6 percent by bicycle, and 2 percent by car.

Table 1: Commuting in Tokyo

Mode	% of Commutes		Average Commute Time (min)	
	Overall	Central Tokyo	Overall	Central Tokyo
Train	54%	70%	63	59
Car	23%	5%	28	36
Cycle	10%	12%	17	17
Foot	6%	8%	16	15
Bus	2%	3%	38	38

model to match the travel times reported in the survey. Specifically, all computed travel times are multiplied by a constant for each commuting-zone pair such that the average travel times computed across origin and destination census-tract pairs within commuter-zone pairs are equated to the average travel times reported in the survey.[13] Whenever there are no reported travel times between two zones (because there are no commuters between the zones), we use predicted values from a regression of reported travels times on computed travel times. Additional details can be found in the appendix.

Municipality-level average earnings (of residents), which are used to identify the Fréchet heterogeneity parameter, is obtained from average taxable earnings information that is published by all muncipalities each year. The same wage data is used to validate our model by comparing the wage data with the wages that were inverted from our model.[14]

Additional information on data sources can be found in the appendix.

5 Model

This section considers a model similar to that of Ahlfeldt et al. 2015. The model exhibits agglomeration forces in productivity and amenities to account for the empirical regularity of agglomeration that characterizes urban areas. Disperion forces in the form of commuting costs and an inelastic supply of land are also incorporated to counteract these forces. The main difference relative to Ahlfeldt et al. 2015 is that their model allows amenity and productivity externalities to spillover to neighboring locations, whereas we follow Tsivanidis 2023 and Allen and Arkolakis 2014 and model spillovers to be restricted to a particular location. This is motivated both by the empirical finding (for example, in Ahlfeldt et al. 2015) that amenity and productivity spillovers decay rapidly across space, and also to simplify estimation. Another difference is that our model allows for multiple modes of transport (trains, cars, bus, bicycle, and foot) in order to capture how commuters change their commuting mode when the train system is removed.

Locations in a city are indexed by $i \in I$ and differ in their primitives by their exogenous levels of amenities $\bar{U}_i$, productivities $\bar{A}_i$, supply of floor space H_i, and the time t_{ijm} it takes to commute between locations using mode m. A continuum of workers with mass $\bar{L}$ who are perfectly mobile

[13] This approach follows Miyauchi, Nakajima, and S. Redding 2021.

[14] See additional discussion in section 7.2.

choose where to live and work and have Cobb-Douglas preferences over a freely traded numéraire good and housing (i.e. floorspace) while longer commutes are modeled to decrease utility. They decide whether or not to move to the city before observing idiosyncratic utility shocks for each possible combination of residence locations, workplace locations, and mode of transport, and then choose their residence, workplace, and travel mode to maximize their utility. Firms produce a freely traded final good with a Cobb-Douglas technology using labor and commercial floor space as inputs.[15] Floorspace is allocated efficiently between residential and commercial land uses.

Indirect utility for an individual living in location i, working in location j and commuting on mode m is given by:

$$V_{ijm}(\omega) = \frac{U_i w_j r_{Ri}^{\beta-1} \epsilon_{ijm}(\omega)}{d_{ijm}} \tag{1}$$

where ϵ_{ijm} is an idiosyncratic preference shock for worker ω commuting between residence i and workplace j using mode m, β is the expenditure on the consumption good, and U_i is the level of amenities. We introduce spillovers to amenities, U_i, so that it is determined endogenously by the number of residents and the exogenous amenity level

$$U_i = \bar{U}_i L_{Ri}^{\mu_U}, \tag{2}$$

where L_{Ri} is the supply of residents and μ_U is the elasticity of amenities with respect to residential population. Commuting costs increase exponentially according to:

$$d_{ijm} = \exp(\kappa_{0m} + \kappa_{1m} t_{ijm}) \tag{3}$$

where κ_{1m} governs the elasticity of commuting costs with respect to time for mode m and κ_{0m} is a mode-specific preference shifter that is independent of travel times.[16] The idiosyncratic preference shocks ϵ_{ijm} are drawn from an i.i.d. Fréchet distribution with shape parameter θ and unit scale:

$$F(\epsilon_{ijm}) = e^{-\epsilon_{ijm}^{-\theta}}, \theta > 1, \epsilon_{ijm} > 0, \tag{4}$$

which implies that the number of commuters for each residence-workplace-mode combination is given by:

$$L_{ijm} = \bar{L} \left(\frac{\gamma}{\bar{V}}\right)^{\theta} \left(\frac{U_i w_j r_{Ri}^{\beta-1}}{d_{ijm}}\right)^{\theta} \tag{5}$$

where $\bar{L}$, the mass of workers, will be set to equal the total number of workers in the Tokyo metro,

[15] As discussed in Ahlfeldt et al. 2015 and modeled in Tsivanidis 2023, this model is isomorphic to one where idiosyncratic shocks and commuting costs affect the effective units of labor, rather than the utility of workers directly. The reason for this isomorphism is that both the idiosyncratic shock and commuting costs enter indrect utility multiplicatively with wages.

[16] We assess the validity of the exponential assumption in section C.4

and $\bar{V}$ is the average indirect utility:

$$\bar{V} = \gamma \left[\sum_j \sum_i \sum_m \left(\frac{U_i w_j r_{Ri}^{\beta-1}}{d_{ijm}} \right)^{\theta} \right]^{1/\theta} \tag{6}$$

and $\gamma \equiv \Gamma\left(\frac{\theta-1}{\theta}\right)$ is a constant that is implied by the Fréchet distribution. $\Gamma(\cdot)$ is the Gamma function. The above expression for average utility is implied by integrating the individual indirect utilities over the Fréchet distribution of idiosyncratic preferences. Equation (5) also implies the following expression for the probability that a worker resides in location i and works in location j and commutes using mode m:

$$\lambda_{ijm} \equiv \frac{L_{ijm}}{\sum_{i,j,m} L_{ijm}} = \frac{\left(\frac{U_i w_j r_{Ri}^{\beta-1}}{d_{ijm}} \right)^{\theta}}{\sum_{i,j,m} \left(\frac{U_i w_j r_{Ri}^{\beta-1}}{d_{ijm}} \right)^{\theta}}. \tag{7}$$

According to equation (5), which exhibits a gravity equation in commuting, θ is therefore the elasticity of workers with respect to commuting costs as well as the elasticity with respect to wages and amenities. The reason for this result is that because the Fréchet shape parameter θ determines the heterogeneity of the distribution of idiosyncratic preferences (a higher θ leads to a lower variance in idiosyncratic preference shocks), the parameter dictates the responsiveness of the population to changes in location characterisitcs, and hence the elasticity of the population with respect to location characteristics, including wages.

Defining the overall commuting cost between locations i and j as:

$$d_{ij} \equiv \left[\sum_m d_{ijm}^{-\theta} \right]^{-1/\theta} \tag{8}$$

yields similar expressions for commuter flows and commuting probailities between location i and j:

$$L_{ij} = \bar{L} \left(\frac{\gamma}{\bar{V}} \right)^{\theta} \left(\frac{U_i w_j r_{Ri}^{\beta-1}}{d_{ij}} \right)^{\theta} \tag{9}$$

$$\lambda_{ij} \equiv \sum_m \lambda_{ijm} = \frac{\left(\frac{U_i w_j r_{Ri}^{\beta-1}}{d_{ij}} \right)^{\theta}}{\sum_{i,j} \left(\frac{U_i w_j r_{Ri}^{\beta-1}}{d_{ij}} \right)^{\theta}}. \tag{10}$$

Summing equation (9) across locations yields the supply of residents and the supply of workers to each location:

$$L_{Ri} = \bar{L} \left(\frac{\gamma}{\bar{V}} \right)^{\theta} \left(U_i r_{Ri}^{\beta-1} \right)^{\theta} \Phi_{Ri} \tag{11}$$

$$L_{Fj} = \bar{L}\left(\frac{\gamma}{\bar{V}}\right)^{\theta} w_j^{\theta}\Phi_{Fj} \tag{12}$$

where we refer to $\Phi_{Ri} \equiv \sum_j (w_j/d_{ij})^{\theta}$ as residential commuter market access (RCMA) and $\Phi_{Fj} \equiv \sum_i (U_i r_{Ri}^{\beta-1}/d_{ij})^{\theta}$ as firm commuter market access (FCMA), following Tsivanidis 2023. RCMA summarizes a location i's geographical access to firms, while FCMA summarizes a location j's access to residents. In the gravity literature, these terms are the multilateral resistance terms such that conditional on transportation costs d_ij between two locations i and j, commuter flows between those locations increases if they are more remote, as measured by RCMA and FCMA. Thus, the supply of residents to a location is increasing in the location's amenities and the commuting access to firms and is decreasing in the price of floor space. Meanwhile, the supply of workers to a location is increasing in the wages of the firms in the location and the commuting access to residents.

Commuting market clearing implies that the number of workers in location j is equal to the number of residents commuting to location j from across the city:

$$L_{Fj} = \sum_i L_{ij} = \sum_i L_{Ri}\lambda_{ij|i} \tag{13}$$

where $\lambda_{ij|i}$ is the probability that a worker residing in location i works in location j (i.e. the probability of working in location j, conditional on residing in location i):

$$\lambda_{ij|i} \equiv \frac{L_{ij}}{L_{Ri}} = \frac{(w_j/d_{ij})^{\theta}}{\sum_{j'}(w_{j'}/d_{ij'})^{\theta}} = \frac{(w_j^{\theta}/d_{ij})^{\theta}}{\Phi_{Ri}} \tag{14}$$

Labor market clearing implies that the population of workers across locations must sum to the total population, and likewise for the population of residents: $\sum_i L_{Ri} = \bar{L}$ and $\sum_j L_{Fj} = \bar{L}$.

Residents spend a $1-\beta$ share of their income on housing, and the demand for housing is therefore given by:

$$H_{Ri} = L_{Ri}\bar{y}_i\frac{1-\beta}{r_{Ri}} \tag{15}$$

where $\bar{y}_i \equiv E[w|i]$ is the average income of residents who reside in location i, and is equal to the average of wages across workplaces, weighted by the share of the residents in location i that work in each workplace:

$$\bar{y}_i = \sum_j \lambda_{ij|i} w_j. \tag{16}$$

Firms in each location produce the freely traded good with a Cobb-Douglas technology $Y_j = A_j L_{Fj}^{1-\alpha} H_{Fj}^{\alpha}$, where L_{Fj} is labor, H_{Fj} is commercial floor space, and A_j is the hicks-neutral productivity that incorporates spillovers according to

$$A_j = \bar{A}_j L_{Fj}^{\mu_A}, \tag{17}$$

where μ_A is the elasticity of productivity with respect to employment. Profit maximization and

normalizing the price of the good ($p = 1$) implies the unit cost function:

$$1 = \frac{1}{A_j}\left(\frac{w_j}{\alpha}\right)^{\alpha}\left(\frac{r_{Fj}}{1-\alpha}\right)^{1-\alpha} \tag{18}$$

where r_{Fj} is the price of commercial floor space. Rearranging implies the following expression for the price of floor space:

$$r_{Fj} = (1-\alpha)\left(\frac{\alpha}{w_j}\right)^{\alpha/(1-\alpha)} A_j^{1/(1-\alpha)}, \tag{19}$$

and the firm's first-order conditions implies the following expressions for labor demand and demand for floorspace:

$$L_{Fj} = \left(\frac{\alpha A_j}{w_j}\right)^{1/(1-\alpha)} H_{Fj} \tag{20}$$

$$H_{Fj} = \left(\frac{(1-\alpha)A_j}{r_{Fj}}\right)^{1/\alpha} L_{Fj} \tag{21}$$

The above equations are satisfied whenever there is production in a region(corresponding to when $L_{Fj} > 0$ in our data).

Floor space can be used commercially or residentially, and the endogenous fraction of floor space allocated to commercial and residential use are denoted υ_i and $1 - \upsilon_i$ respectively.

The no-arbitrage condition across alternative land uses implies that floor space is used entirely commercially, used entirely residentially, or used in a mix of both if the commercial price of floor space equals the residential price of floor space. The share of floor space used commercially is given by:

$$\begin{aligned} &\upsilon_i = 1 \text{ if } r_{Fi} > r_{Ri} \\ &\upsilon_i \in [0,1] \text{ if } r_{Fi} = r_{Ri} \\ &\upsilon_i = 0 \text{ if } r_{Fi} < r_{Ri} \end{aligned} \tag{22}$$

The observed price of floor space, R_i is assumed to be the maximum of the commercial and residential price of floor space $R_i = \max(r_{Fi}, r_{Ri})$, such that the observed price of floor space is given by:

$$\begin{aligned} &R_i = r_{Fi}, \quad \upsilon_i = 1 \text{ if } r_{Fi} > r_{Ri} \\ &R_i = r_{Ri}, \quad \upsilon_i \in [0,1] \text{ if } r_{Fi} = r_{Ri} \\ &R_i = r_{Fi}, \quad \upsilon_i = 0 \text{ if } r_{Fi} < r_{Ri} \end{aligned} \tag{23}$$

Wherever the density of development is not observed in the data and for the purpose of the counterfactual, we follow the standard approach used in the urban literature (see Ahlfeldt et al. 2015 and Epple, Gordon, and Sieg 2010) and assume that floor space, $H_i = H_{Ri} + H_{Fi}$, is supplied by a competitive construction sector with a Cobb-Douglas production function that uses immobile

land, D, and mobile capital K as inputs. Let

$$H_i = K_i^{\mu} D_i^{1-\mu}$$

Given fixed land and mobile capital (with an exogenous rental rate), the supply and price of floor space is given by:

$$H_i = \varphi_i D_i^{1-\mu} \tag{24}$$

$$R_i = \chi R_{Li}^{1-\mu}, \tag{25}$$

where R_{Li} is the price of land, χ is a constant, and $\varphi_i = K_i^{\mu}$ is referred to as the density of development because it determines the relationship between floor space and land area.

Residential land market clearing implies that the supply of floor space allocated for residential use $(1 - \upsilon_i)H_i$ is equal to the demand for floor space (H_{Ri}):

$$(1 - \upsilon_i)H_i = \frac{(1-\beta)\bar{y}_i L_{Ri}}{r_{Ri}}, \tag{26}$$

where $\bar{y} = E[w|i]$ is as in equation (16). Commercial land market clearing implies that the supply of floor space allocated for commercial use $\upsilon_i H_i$ is equal to the demand for commercial floor space. The Cobb-douglas production function then implies:

$$\upsilon_j H_j = L_{Fj} \left(\frac{A_i(1-\alpha)}{r_{Fi}} \right)^{1/\alpha} \tag{27}$$

For the case of exogenous productivities and amenities (no spillovers), Ahlfeldt et al. 2015 provide proofs guaranteeing existence of a unique general equilibrium of this model. They also show that there is a unique mapping from the observed variables to unobserved location characteristics including production and residential fundamentals even when we allow for spillovers, as is described in the following proposition, which we state without proof. Defining the vectors: $\mathbf{R} \equiv [R_{Li}]$, $\mathbf{L_R} \equiv [L_{Ri}]$, $\mathbf{L_F} \equiv [L_{Fj}]$, $\mathbf{D} \equiv [D_i]$, and $\mathbf{t} \equiv [t_{ijm}]$, $\kappa_{\mathbf{0}} \equiv [\kappa_{0m}]$, $\kappa_{\mathbf{1}} \equiv [\kappa_{1m}]$, $\mathbf{A} \equiv [A_i]$, and $\mathbf{U} \equiv [U_i]$, we then have:

Proposition 1. *Given known values for the parameters* $\{\alpha, \beta, \mu, \theta, \kappa_{\mathbf{0}}, \kappa_{\mathbf{1}}, \mu_U, \mu_A\}$ *and the observed data* $\{\mathbf{R}, \mathbf{L_R}, \mathbf{L_F}, \mathbf{t}\}$*, then there exists unique vectors of unobserved location characteristics* $\{\mathbf{A}, \mathbf{U}\}$ *that are consistent with the data being an equilibrium of the model.*

This proposition guarantees that with our data, we are able to recover location fundamentals, which we can then use to compute counterfactuals and perform welfare analysis. Interested readers are referred to the supplementary appendix of Ahlfeldt et al. 2015 where the authors prove the proposition in steps using Lemmas S.6-S.11. The differences between our proposition and theirs is as follows. First, the present paper defines multiple modes (thus, commute costs are $d_{ij} = [\sum_m \exp(-\theta(\kappa_{0m} + \kappa_{1m} t_{ijm}))]^{-1/\theta}$) whereas Ahlfeldt et al. 2015 use a single mode (with commute

costs: $d_{ij} = exp(\kappa t_{ij})$). Second, the present paper uses a different choice of normalization.[17] Last, the present paper does not recover the density of development φ_i nor requires data on land area, as these variables are not necessary for the counterfactuals we perform.[18]

In summary, we have presented a model that can be summarized by the following system of equations (reproduced from equations (6) (average utility), (11) (supply of residents), (12) (supply of workers), (26) (residential land-market clearing), (27) (commercial land-market clearing), (19) (firm cost minimization), (16) (worker income), (22) (floor-space no-arbitrage condition), (8) (commuting-cost aggregation), (2) (amenity spillovers), (17) (production spillovers), and the expressions for RCMA and FCMA):

$$\bar{V} = \gamma \left[\sum_j \sum_i \sum_m \left(\frac{U_i w_j r_{Ri}^{\beta-1}}{d_{ijm}} \right)^{\theta} \right]^{1/\theta}$$

$$L_{Ri} = \bar{L} \left(\frac{\gamma}{\bar{V}} \right)^{\theta} \left(U_i r_{Ri}^{\beta-1} \right)^{\theta} \Phi_{Ri}$$

$$L_{Fj} = \bar{L} \left(\frac{\gamma}{\bar{V}} \right)^{\theta} w_j^{\theta} \Phi_{Fj}$$

$$(1 - \upsilon_i) H_i = \frac{(1-\beta)\bar{y}_i L_{Ri}}{r_{Ri}}$$

$$\upsilon_j H_j = L_{Fj} \left(\frac{A_i (1-\alpha)}{r_{Fi}} \right)^{1/\alpha}$$

$$\upsilon_i = 1 \text{ if } r_{Fi} > r_{Ri}$$

$$\upsilon_i \in [0,1] \text{ if } r_{Fi} = r_{Ri}$$

$$\upsilon_i = 0 \text{ if } r_{Fi} < r_{Ri}$$

$$r_{Fj} = (1-\alpha) \left(\frac{\alpha}{w_j} \right)^{\alpha/(1-\alpha)} A_j^{1/(1-\alpha)}$$

$$\bar{y}_i = \sum_j \lambda_{ij|i} w_j$$

$$\lambda_{ij|i} = \frac{(w_j / d_{ij})^{\theta}}{\sum_{j'} (w_{j'} / d_{ij'})^{\theta}}$$

$$U_i = \bar{U}_i L_{Ri}^{\mu_U}$$

$$A_j = \bar{A}_j L_{Fj}^{\mu_A}$$

$$\Phi_{Ri} \equiv \sum_j (w_j / d_{ij})^{\theta}$$

[17] Specifically, Ahlfeldt et al. 2015 define a set of adjusted variables in Section S.3 that the authors prove are mapped uniquely from the data. However, those adjusted variables are precisely our variables $\{\mathbf{A}, \mathbf{U}, \varphi\}$ given the assumptions we have made in our model. More specificially, if we use the notation from Ahlfeldt et al. 2015, our setting sets $E_i = T_i = \xi_i = \zeta_{Ri} = 1$, and then the adjusted variables are equal to the unadjusted variables (which we use), and therefore their results translate directly to our setting.

[18] Ahlfeldt et al. 2015 show that the model without spillovers (i.e. $\mu_U = \mu_A = 0$) has a unique equilibrium. Allen, Arkolakis, and Li 2024 provide sufficient and necessary conditions for uniqueness and sufficient and necessary conditions for multiplicity for a similar model where land is separated into residential-specific and commercial-specific land.

$$\Phi_{Fj} \equiv \sum_i (U_i r_{Ri}^{\beta-1}/d_{ij})^\theta$$

$$d_{ij} = \left[\sum_m \exp(-\theta(\kappa_{0m} + \kappa_{1m} t_{ijm}))\right]^{-1/\theta}$$

6 Model with Road Congestion

We model road congestion using the routing-choice framework from Allen and Arkolakis 2022. In addition to the choice of residence, i, workplace j, and transportation mode m, workers also choose a route r. Indirect utility is given by:

$$V_{ijmr}(\omega) = \frac{U_i w_j r_{Ri}^{\beta-1} \epsilon_{ijmr}(\omega)}{d_{ijmr}} \tag{28}$$

where we have replaced commuting costs d_{ijm} with d_{ijmr} and idiosyncratic preference shock ϵ_{ijm} with ϵ_{ijmr}. Agents choose their workplace, residence, and transportation mode first, and then choose their route in a nested structure. In other words, preference shocks are drawn from a nested Fréchet distribution with shape parameter θ in the upper nest and parameter ρ in the lower nest. The joint cumuluative distribution function is from the generalized extreme-value family and is given by:

$$F(\{\epsilon_{ijmr}\}) = \exp\left(-\sum_{i,j,m}\left(\sum_{r\in\mathscr{R}_{ijm}} \epsilon_{ijmr}^{-\rho}\right)^{\frac{-\theta}{\rho}}\right), \theta > 1, \theta < \rho, \epsilon_{ijmr} > 0, \tag{29}$$

where $\mathscr{R}_{ijm}$ is the set of routes between residence i and workplace j using mode m.

We focus our attention on automobile routing as other modes (bicycle, walking, bus) are much less suceptible to congestion or are not used nearly as much as cars.

Thus, we assume that there is only one route for modes other than cars and therefore the commuting cost for these modes is the same as before: $d_{ijmr} = d_{ijm}$ for $m \in \{\text{walk, cycle, bus}\}$. For cars, commuting costs between locations i and j are implied by the nested Fréchet structure:

$$d_{ijm} = \left[\sum_{r\in\mathscr{R}_{ijm}} d_{ijmr}^{-\rho}\right]^{\frac{-1}{\rho}} \tag{30}$$

and implies the same indirect utility as before (in equation (28)), but with the commuting cost for cars, d_{ijm}, given by equation (30). Each location (in our case, a census tract) is a node on the network and is connected to other nodes either directly (if there is a direct link between the two nodes) or indirectly (via a series of other nodes). Associated with each link is a travel time, and a route is defined by a sequence of links such that the total travel time is the sum of the travel times of each link in the route.

We denote the commuting costs for cars between two nodes on a direct link as d^L_{kl} where k is the origin node and l is the destination node.[19] The commuting costs for cars between two nodes taking a route r is then given by:

$$d_r = \prod_{(k,l)\in r} d^L_{kl} \tag{31}$$

where the product is taken over all links in the route r. Analagous to the previous section, we parametrize link-level costs as $d^L_{kl} = \exp(\kappa_{1,\text{car}} t_{kl})$, such that route-level costs are given by:

$$d_r = \exp\left(\kappa_{1,\text{car}} \sum_{(k,l)\in r} t_{kl}\right). \tag{32}$$

It proves conveneient to describe the network using the matrix:

$$\mathbf{A} \equiv \left[\left(d^L_{kl}\right)^{-\rho}\right] \tag{33}$$

where $d^L_{kl} = \exp(\kappa_{1,\text{car}} t_{kl})$ whenever the nodes k and l are directly connected, and $d^L_{kl} = \infty$ otherwise. "Self-loops," links that start and end in the same node, are forbidden: $d^L_{kk} = \infty$, such that a commuter must leave the node once and then return, in order to commute in the same workplace as their residence. Allen and Arkolakis 2023 show that the Fréchet distribution implies that commuting costs for cars (Equation 30) can be expressed in terms of the Leontief inverse of the matrix $\mathbf{A}$. Denoting $\mathbf{B} \equiv \left[d^{-\rho}_{ijm}\right]$, they show:

$$\mathbf{B} = (\mathbf{I} - \mathbf{A})^{-1} \tag{34}$$

where $\mathbf{I}$ is the identity matrix. The commuting costs for cars between two nodes i and j is then given by $d_{ijm} = B_{ij}^{-\frac{1}{\rho}}$ where B_{ij} are the (i,j)-th element of matrix $\mathbf{B}$, where once again, we have suppressed the index m as it is assumed that we are only considering car commutes.

In addition, the authors show that *link intensity*, the expected number of times that a link k,l is used for commutes between an origin i and destination j, can be expressed as

$$\pi^{kl}_{ij} = \left(\frac{d_{ijm}}{d_{ikm} d^L_{kl} d_{ljm}}\right)^{\rho} \tag{35}$$

Traffic flows through link k,l are then the sum across all pairs of nodes i,j of the link intensity, weighted by the number of commuters between i and j:

$$Q_{kl} = \sum_{i,j} \pi^{kl}_{ij} L_{ij\text{car}}. \tag{36}$$

Travel times between link kl are parameterized to be a sum of an exogenous component of travel

[19] We suppress the i,j,m subscript for notational clarity. m = car and the link-level cost is independent of the the trip origin i or trip destination j.

time, $\bar{t}_{kl}$, and an endogenous component that is a function of traffic flows, $t' = G_{kl}Q_{kl}^{\sigma}$, such that the total time passing through the link is:

$$t_{kl} = \bar{t}_{kl} + G_{kl}Q_{kl}^{\sigma}. \tag{37}$$

The exogenous component, $\bar{t}_{kl}$, represents the travel time on the link if there was no traffic congestion, the parameter σ is the elasticity of the endogenous component of travel times with respect to traffic flows, and G_{kl} is a link-specific constant that captures the width of roads and other link-level attributes that mediate the impact of traffic on travel times. Allen and Arkolakis 2023 provide sufficient conditions for the existence of equilibria and sufficient conditions for a unique equilibrium for a similar model with fixed commercial and residential land.

7 Estimation

The necessary data for our analysis are $\{\mathbf{R_L}, \mathbf{L_R}, \mathbf{L_F}, \mathbf{t}\}$, which are the price of land, residential population, employment, land area, and travel times.[20] Employment and population numbers are obtained directly from Census data. Land prices are obtained from the Official Land Price Announcements. Travel times are calculated using a combination of the street map (using Openroute Services), the geography of the train network (taken from MLIT), and reported travel times from the Tokyo Person-Trip Survey, as discussed in section (B).

Converting land prices to floor prices using equation (25) then yields the variables required in Proposition 1: $\{\mathbf{R}, \mathbf{L_R}, \mathbf{L_F}, \mathbf{t}\}$.

We then proceed as follows in order to estimate the necessary parameters of the model. First, we use the gravity structure of the model to estimate the semi-elasticity of commuting flows with respect to travel time for each mode, corresponding to the product, $\theta\kappa_{0m}$. Second, we use this elasticity together with employment and population data, and bilateral travel times for each mode to recover transformed wages, w_i^{θ}, which are recoverable because they are proportional to commuter flows, conditional on commuting costs and multilateral resistance (RCMA). We match the model's prediction of the dispersion of log earnings across locations with that of Tokyo to separately recover θ and $\kappa_{\mathbf{0}}$, and thus wages. Intuitvely, a higher θ implies less heterogeneity in preference shocks, higher sensitivity to wages, and thus lower dispersion in wages conditional on observed commuter flows. Next, we use the wages, together with floor-space prices, residential population, and travel times to back out productivity and amenity levels for every location. Last, we use an instrumental variables strategy to estimate the spillover parameters μ_U and μ_A that correspond to the elasticity of amenities with respect to residential population and the elasticity of productivity with respect to employment, respectively. The instrumental-variables estimation is ongoing and so details of the estimation are relegated to the appendix.

[20] Land area, $\mathbf{D}$, while necessary for computing the density of floorspace, is not necessary for other parts of the analysis, since it is assumed to be constant in counterfactuals.

7.1 Parameters from the Urban literature

Of the parameters in our model $\{\alpha, \beta, \mu, \theta, \kappa, \mu_U, \mu_A\}$, a subset will be taken from previous estimates in the literature. We set the share of consumer expenditure on residential floor space $(1-\beta)$ equal to 0.25, consistent with the estimates from Davis and Ortalo-Magné 2011. The parameter that governs the share of firm expenditure on commercial floor space $(1-\alpha)$ is set to 0.2, consistent with findings from Valentinyi and Herrendorf 2008. Last, the share of land in construction costs, $(1-\mu)$ is set to 0.25, consistent with Combes, Duranton, and Gobillon 2019 and Epple, Gordon, and Sieg 2010. For the model with congestion, we set the elasticity of the endogenous component of travel times with respect to traffic flows, σ, to 0.35, which is about halfway between the two IV estimates from Allen and Arkolakis 2023 between 0.1 and 0.5. Unlike their estimate which is the elasticity of travel times with respect to traffic flows, our parameter is the elasticity of the endogenous component of travel times with respect to traffic flows. It would be appropriate to set our parameter estimate higher to account for these differences. However, given the substantial uncertainty in the literature(e.g. the standard error for the preferred estimate in Allen and Arkolakis 2022 of 0.49 is 0.28), we choose the estimate of 0.35 for now, and plan to perform sensitivity analyses in the future.

7.2 Estimating parameters governing commuting costs

First we use the commuter flow data from the Personal Trip survey to estimate a gravity equation of the following form:

$$\ln \pi_{ijm} = \delta_m + \gamma_i + \vartheta_j - \theta\kappa t_{ijm} \tag{38}$$

where π_{ijm} is the share of workers that have selected residence in location i and workplace in location j, as expressed in equation (7).[21]

Given travel times and observed variables, our model predicts a perfect fit of the commuting flow data according to equation (5). However, given the observed commuting flows and travel times are self-reported and from a finite sample from a finite population, measurement error is likely to be present. There is potential misspecification in functional form and potential heterogeneity across individuals. We therefore add an error term ε_{ijm} yielding the regression model:

$$\ln \pi_{ijm} = \delta_m + \gamma_i + \vartheta_j - \theta\kappa t_{ijm} + \varepsilon_{ijm} \tag{39}$$

where we assume that the error ε_{ijm} is uncorrelated with travel times t_{ijm}.

In addition, our model predicts commuting flows at the census tract level, whereas the commuter-flows data is at the commuting-district level, whose population is about 7-8 times larger than that of a census tract on average. We therefore require additional assumptions to take our model to the data. We follow the strategy in Ahlfeldt et al. 2015 and first take means across pairs of blocks to

[21] We use different notation for the model-implied probabilities, λ_{ijm}, and the observed probabilities, π_{ijm}, to emphasize the potential difference between the two quantities in other models. See for example, Dingel and Tintelnot 2023. In our case, we however assume that $\pi_{ijm} = \lambda_{ijm}$

obtain

$$\ln \pi_{IJm} = \delta_m + \gamma_I + \vartheta_J - \theta\kappa t_{IJm} + \varepsilon_{IJm} \tag{40}$$

where I and J are commuting districts, t_{IJm} is the average travel time between the two commuting districts across the census-tract pairs in those districts, and π_{IJm} is the approximated average of log commuting probabilities across all pairs of tracts in commuting districts I and J (i.e. the average of $\ln \pi_{ijm}$). We approximate the average of $\ln \pi_{ijm}$ using the observed log district commuting probabilities, and the average of the travel times across census tracts t_{IJm} with the average travel time across commuters, which is given in the survey data.[22] In summary, we run a Poisson Pseudo-Maximum Likelihood (PPML) estimation on commuting flows from our Personal-Trip-Survey data on origin-district fixed effects, destination-district fixed effects, mode fixed effects, and on the average travel time between commuting districts.The estimated coefficient on travel time is the semi-elasticity of commuting flows with respect to travel time for a particular mode and is equal to the product of the preference heterogeneity parameter θ and the commuting costs parameter κ_{1m}. Mode-specific constants δ_m are equal to product of the heterogeneity parameter and mode-specific fixed commuting costs ($\theta \cdot \kappa_{0m}$). In practice, our gravity estimates from the above methodology led to a poor fit with observed commuting shares (for example, the share of commuting by train was about 30 percent when the commuter-flow data showed it to be at 53 percent). We believe this discrepency arose due to failure of the large of large numbers to hold, resulting from a finite population size, a modest sampling rate(about 2%), and small commuting districts (about 15 thousand).[23] Indeed, only 1.6 percent of possible origin-destination-mode combinations had non-zero trips at this level of aggregation. We therefore aggregate the data to a higher level of aggregation, where the population of each region is about 750 thousand on average, and repeat the PPML exercise.
[24]

Given estimates of $\theta\kappa_{0m}$ and $\theta\kappa_{1m}$, we use the labor market clearing condition (equation (13)) to obtain:

$$L_{Fj} = \sum_i L_{Ri}\lambda_{ij|i} = \sum_i L_{Ri}\frac{(w_j/d_{ij})^\theta}{\sum_{j'}(w_{j'}/d_{ij'})^\theta} \tag{41}$$

Defining transformed variables $\tilde{w}_i = w_i^\theta$ and $\tilde{d}_{ij} = d_{ij}^\theta = [\sum_m \exp(-\theta\kappa_{0m} - \theta\kappa_{1m}t_{ijm}))]^{-1}$, we rewrite the above equation as:

$$L_{Fj} = \sum_i L_{Ri}\frac{\tilde{w}_j/\tilde{d}_{ij}}{\sum_{j'}\tilde{w}_{j'}/\tilde{d}_{ij'}} \tag{42}$$

Data on travel times, employment, and population data are then used to recover transformed

[22]The approximation for the average of $\ln \pi_{ijm}$ is based on taking the log of the averages (the commuter-flow data), rather than the log of the averages (the true commuting probability), and therefore the discrepancy depends on the extent to which Jensen's inequality leads to a discrepency betwen the average of the logs and the log of the average. In practice, Ahlfeldt et al. 2015 find that the approximation is close in the sense that it leads to a negligible difference in parameter estimates.

[23]See Dingel and Tintelnot 2023 for a discussion of the challenges in calibrating transport costs in granular settings.

[24]We are currently undergoing various checks to see if our estimates our robust to various geographical definitions and data specifications.

wages for the N locations using these N equations (up to a choice of normalization).[25] This in turn allows us to compute the conditional commuting probabilities ($\lambda_{ij|i}$).

Given transformed commute costs $\tilde{d}_{ij}$ and transformed wages ($\tilde{w}_i$) we calibrate the heterogeneity parameter θ by using our data on the dispersion of log earnings by residence location across the districts of Tokyo. Because the Fréchet parameter applies a monotonic transformation to wages which adjusts the inequality in wages across locations, we use data on the dispersion of log earnings across districts, which are a weighted average of workplace wages, to calibrate θ. In particular, we use municipality-level data on average taxable earnings and set θ to equate the variance of log average earnings in the model (the population-weighted average of $\bar{y}_i \equiv E[w|i]$ from equation (16) within a municipality) with that of the taxable earnings data. [26] Once the Fréchet parameter is estimated, then together with our estimates of $\theta\kappa_{0m}$ and $\theta\kappa_{1m}$, we recover the semi-elasticity of commute costs with respect to travel time, κ_{1m} for each mode and mode-specific fixed costs, κ_{0m}.[27]

7.3 Recovering Productivity, Amenities, and Density of Development

We continue to use the gravity structure of the model to recover productivities, amenities, and density of development. The unit cost function (equation (18)) implies:

$$\ln A_j = (1-\alpha)\ln r_{Fj} + \alpha \ln w_j - \alpha \ln \alpha - (1-\alpha)\ln(1-\alpha) \tag{43}$$

Given values for the price of floor space and wages, we recover productivities up to a normalization. For calibrating the model, we set the geometric mean of floor-prices to 1, and the geometric mean of wages to 1, which then implies values for productivities. To recover amenities, we use the supply of residents in each location (equation (11)) and divide by its geometric mean to obtain the following expression for amenities:

$$\ln \frac{U_i}{\bar{U}} = \frac{1}{\theta}\ln \frac{L_{Ri}}{\bar{L}_R} + (1-\beta)\ln \frac{r_{Ri}}{\bar{r}_R} - \frac{1}{\theta}\ln \frac{\Phi_{Ri}}{\bar{\Phi}_R} \tag{44}$$

where $\Phi_{Ri} = \sum_j (w_j/d_{ij})^\theta$ as defined previously and an overbar ($\bar{\cdot}$) refers to the geometric mean across census tracts. Note that the reservation utility has been eliminated from the above equation because it is common across locations. We then set the left hand side as our amenities respectively (i.e. $\ln U_i = \frac{1}{\theta}\ln \frac{L_{Ri}}{\bar{L}_R} + (1-\beta)\ln \frac{r_{Ri}}{\bar{r}_R} - \frac{1}{\theta}\ln \frac{\Phi_{Ri}}{\bar{\Phi}_R}$). By dividing the supply-of-residents equation by the geometric mean, we are setting a normalization for amenities such that the geometric means of amenities is equal to 1.

[25] When the unit of analysis for the commuter-flows regression is the same as the rest of the analysis, an alternative approach is to utilize the destination fixed effects from the commuter-flows gravity regression, as these fixed effects are equal to log transformed wages.

[26] Ahlfeldt et al. 2015 instead match the variance of log workplace wages between the model and the Berlin data and obtain an estimate of between 6.5 to 6.8. We are not aware of reliable workplace wage data.

[27] Ahlfeldt et al. 2015 obtain a value of $\kappa = 0.01$. Tsivanidis 2023 obtain a similar value of 0.012 by estimating a discerte choice model of commuting mode and destination choice using a mobility survey of Bogota, where κ is identified from the sensitivity of individuals' mode choices to differences in times across modes within particular commutes.

Combining commercial and residential land market clearing conditions (equations (27) and (26)) yields the following expression involving the density of development:

$$H_i = \varphi_i D_i^{1-\mu} = (1-\upsilon_i)H_i + \upsilon_i H_i = \frac{(1-\beta)\bar{y}_i L_{Ri}}{r_{Ri}} + L_{Fj}\left(\frac{A_i(1-\alpha)}{r_{Fi}}\right)^{1/\alpha} \tag{45}$$

where, as defined previously, $\bar{y}_i = \sum_j \lambda_{ij|i} w_j$, and we have that $r_{Ri} = r_{Fi}$ wherever both commercial and residential floor space are used in location i. Thus, we can use our values of parameters for $\{\alpha, \beta, \mu, \theta, \kappa\}$ together with inputs of wages, land prices, workplace population, residence population, travel times to solve for productivities, amenities, and density of development up to a normalization.[28]

7.4 Estimation of Spillovers

Between 1995 and 2020, ∼350km of new track was added to Tokyo's railway network, about a 7 percent increase in track length. We treat the change in employment and residential population induced by the change in commuting accessility as plausibly exogenous and estimate its impact on productivity and amenities, respectively. Because the estimation is currently ongoing, we relegate the details to the appendix. For the quantification, we utilize spillover estimates from the recent urban literature, which we present in Table 4.

7.5 Calibrating the Road Congestion Model

For this draft, we discuss how we implement the calibration at the municipality level. The census-tract-level analysis is ongoing. We use the reported and computed car-travel times together with the previously computed commuting costs to calibrate the congestion model. We use GIS software to identify each location's adjacent locations and travel times from the survey. We believe the surveyed time is typically an overestimate of the link-level travel time, because there are likely fixed costs at the beginning and the end of a trip to enter and exit the network. Thus, we assume that the reported times should be scaled by a fixed factor such that the link-level costs are given by:

$$d_{kl}^L = \exp(\kappa_{1,\text{car}}(t_{kl}^{rep} - \tilde{t})) \tag{46}$$

where t_{kl}^{rep} is the reported travel time between locations k and l, and $\tilde{t}$ is the fixed travel time to get on and off the road network. For hypothetical $\tilde{t}$, commuting-cost matrix can be constructed using equation (34). We then choose $\tilde{t}$ to minimized the sum of squared differences between the

[28] The density of development is exogenous in this model, and therefore housing supply does not adjust endogenously. However, because the density of development is a structural residual, floor space is allowed to adjust between the pre and post period for the purpose of estimating the parameters of the model. In counterfactuals however, the current version of the model keeps the level of floorspace fixed. In future work, we plan to allow floorspace to adjust.

conditional commuting-cost shares in the congestion model and the model without congestion:

$$\min_{\tilde{t}} \sum_{i,j} \left(\frac{d_{Cij}^{-\theta}(\tilde{t})}{\sum_j d_{Cij}^{-\theta}(\tilde{t})} - \frac{d_{ij}^{-\theta}}{\sum_j d_{ij}^{-\theta}} \right)^2 \tag{47}$$

where d_{Cij} denotes the commuting cost between locations i and j in the model with congestion while d_{ij} denotes the commuting cost between locations i and j in the previous model that holds congestion fixed. Given an estimate of $\tilde{t}$ and implied link-level costs d_{kl}^L, link intensity is calculated using equation (35), and traffic flows using equation (36). We view the OpenRouteService travel time as the time that excludes traffic congestion t_{kl}^{OR}, because the service advertises their travel times in this way. However, this travel time also contains the fixed cost. We therefore scale the OpenRouteService travel times by the same time-factor that the reported travel time is scaled, so that

$$\bar{t}_{kl} = t_{kl}^{OR} \frac{t_{kl}}{t_{kl}^{rep}} \tag{48}$$

where $\bar{t}_{kl}$ is the exogenous component of travel time, as discussed in equation (37), and t_{kl} is the link-level travel time that was recovered from the minimum distance estimation. Given these variables we can rearrange equation (37) to recover the link-specific constant G_{kl}, and our model is calibrated.

8 Quantification

Equipped with parameters that have either been estimated or taken from the literature and structural residuals that have been calculated using model-inversion, we then perform two counterfactual simulations. In the first counterfactual, we remove the train system while keeping travel times for other modes of transportation fixed. In the second counterfactual, we allow travel times for cars to change as commuters switching to cars lead to road congestion that increase travel times. Road infrastructure is however held fixed. In future work, we seek to endogenize the density of roads such that road infrastructure can adjust in response to the increased demand for car travel. Parameters from the literature that are used for the quantification are shown in Table 4.

9 Results

Gravity results using OLS and PPML are presented in table 2. Predicted commuter flows implied from PPML specification with mode-specific constants and mode-specific slopes are shown in Figure 3. Given these estimates, we separately identify κ_{0m}, κ_{1m} and θ by matching the model's predictions of the variance of log earnings with that of the data. We estimate the heterogeneity parameter as $\theta = 2.5$ and the implied commuting-cost parameters are shown in Table 3. The results show substantial heterogeneity both in the value of time across different modes and in fixed mode-specific costs, independent of travel time. Trains and buses have the cheapest marginal costs with respect

to a unit of travel time whereas cars cost about 60 percent more than trains per unit of time. The difference may reflect not only differences in amenity value (e.g. a commuter may choose to read on the train but will have to settle for an audiobook if they choose to drive) but also monetary costs (e.g. of fuel). On the other hand, trains have higher fixed costs than walking, cycling or driving. A ten-minute increase in round-trip commuting time on a train is equivalent to a 11 percent increase decrease in wages while a ten-minute increase in round-trip commuting time on a car is equivalent to a 17 percent decrease in wages.

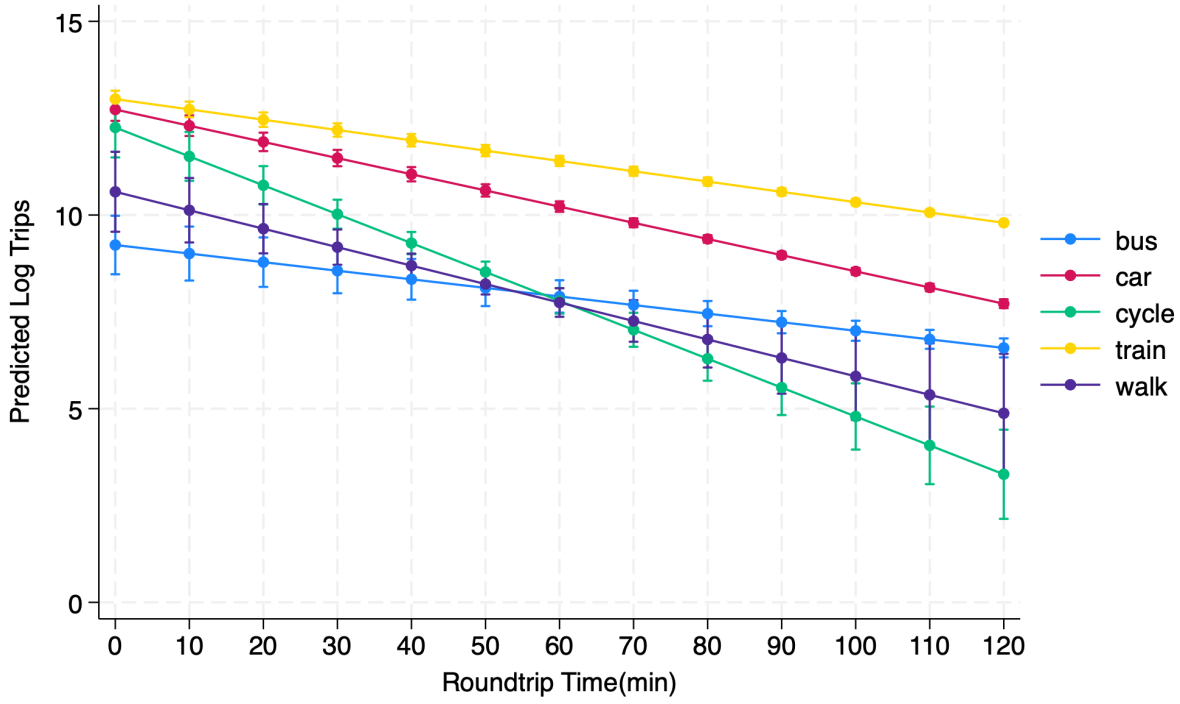


Figure 3: Predicted Commuter Flows from Gravity Model

Predicted commuter flows from the "PPML" "Modes" specification from Table 2. Round-trip travel times are varied from 0-120 minutes for each mode and the predicted log commuter flows are computed for each observation (origin-destination-mode) using the estimated coefficients and origin and destination fixed effects. The figure then reports the average predicted log commuter flows for each mode across these observations. Predicted flows were computed using Stata's "margins" and "marginsplot" commands after estimation with the "ppmlhdfe" command.

The results in Table 3 and our estimate of $\theta = 2.5$ are broadly consistent with previous estimates from the literature. The estimated commuting cost parameters κ_{1m} for trains and buses are similar to those found in Ahlfeldt et al. 2015 and Tsivanidis 2022 of roughly 0.01. However, we estimate separate parameters for different modes and find substantial heterogeneity. The heterogeneity helps to explain why commuting times are signficantly larger for train commuters than car-commuters and also for train-centered cities like Tokyo (about 60 minutes) than in car-centered cities like Los Angeles (about 30 min). The heterogeneity also has implications for welfare and the impact of train removal, as residents will switch to shorter commutes leading to less separation between workplace

Table 2: Gravity Estimates

	OLS Stacked	PPML Stacked	OLS Modes	PPML Modes	OLS Same Slope	PPML Same Slope
Time (min)	-0.017	-0.033			-0.017	-0.033
	(0.001)	(0.001)			(0.001)	(0.001)
Mode # Time (min)						
bus			-0.014	-0.022		
			(0.002)	(0.003)		
car			-0.015	-0.042		
			(0.001)	(0.002)		
cycle			-0.018	-0.075		
			(0.005)	(0.008)		
train			-0.020	-0.027		
			(0.001)	(0.001)		
walk			-0.029	-0.048		
			(0.006)	(0.011)		
Mode						
car			1.244	3.498	1.315	2.077
			(0.269)	(0.412)	(0.105)	(0.213)
cycle			0.328	3.033	-0.015	0.576
			(0.426)	(0.548)	(0.144)	(0.260)
train			3.168	3.768	2.346	3.767
			(0.273)	(0.400)	(0.110)	(0.208)
walk			0.752	1.374	-0.162	0.093
			(0.429)	(0.659)	(0.178)	(0.282)
N	3480	13786	3540	14045	3540	14045
Origin FE			Yes	Yes	Yes	Yes
Destination FE			Yes	Yes	Yes	Yes
Origin-Mode FE	Yes	Yes				
Destination-Mode FE	Yes	Yes				

OLS and PPML estimates of commuter flows on travel times. Travel times refer to average round-trip travel times. "Stacked" estimates estimate one slope coefficient for all modes. "Modes" specification estimate separate slopes for each mode. "Same Slope" allows for mode-specific constants but one slope across all modes. Slope coefficients are interpreted as the log-point change in commuter flows associated with a one-minute increase in commuting time for particular mode, holding constant origin and destination fixed effects (multilateral resistance). Mode-specific constants are interpreted as the log-difference in commuter flows between a mode and bus, when travel times are zero and holding fixed origin and destination fixed effects. The unit of observation is the "big" zone, which contains on average 250 thousand residents. Data on commuter trips and reported times are obtained by Tokyo's Person-Trip Survey. Reported travel times are only available for observations with trips, so travel times for observations with no trips are imputed using predictions from travel times that are computed using the street map and the train network. Details are found in section B.3 of the appendix. Robust standard errors in parentheses. OLS estimates were estimated with Stata's "reghdfe" package while PPML estimates were estimated with Stata's "ppmlhdfe" command. See Correia, Guimarães, and Zylkin 2020 and Correia 2023.

Table 3: Commuting-Cost Parameter Estimates

Mode	Constant (κ_{0m})	Slope (κ_{1m})
Train	-1.5	0.011
Car	-1.4	0.017
Walk	-0.6	0.019
Cycle	-1.2	0.030
Bus	0	0.009

Table 4: Parameters Used in Counterfactual

Parameter	Value	Source
Heterogeneity Parameter (θ)	2.5	Estimated
Cost share of Labor (α)	0.8	Valentinyi and Herrendorf 2008
Expenditure share of final good (β)	0.75	Davis and Ortalo-Magné 2011
Construction sector cost share of physical capital (μ)	0.75	Combes, Duranton, and Gobillon 2019
Production spillover (μ_A)	0.07	Ahlfeldt et al. 2015
Amenity spillovers (μ_U)	0.15	Ahlfeldt et al. 2015

and residence locations and more population dispersal. Introducing road congestion into the model will further magnify these mechanisms.

Our esimate of the Fréchet parameter is similar to those estimated in Tsivanidis 2023 of between 2.7 and 3.3 and lower than the estimate of 6.5 to 6.8 in Ahlfeldt et al. 2015 which use a similar estimation strategy of matching the variance of log wages. The difference may be because incorporating more than one mode requires more preference heterogeeniety to rationalize the data.

Our estimate of the fixed travel time to enter the street network is about 6min. While we are not aware of comparable studies, we believe this figure is plausible given the time it takes to drive to the main road or high way from a residence.

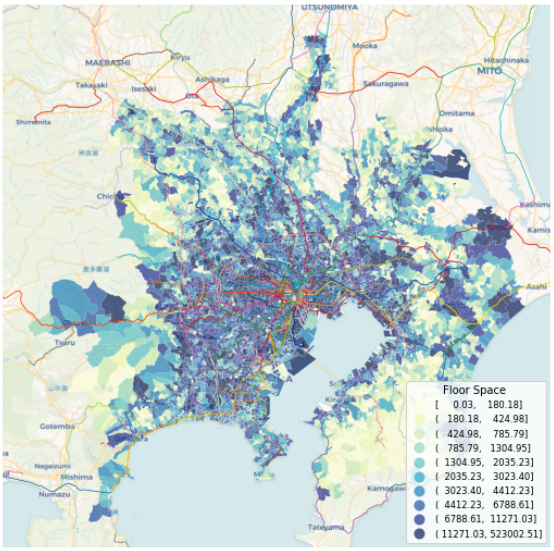
9.1 2020 Equilibrium

The equilibrium variables and fundamentals from the 2020 equilibrium are plotted in Figure 5 and 4. In the 2020 equilibrium, about 30% of the residential population, 47% of employment, and 50% of total land value is contained within 20km of the Imperial Palace (a region that contains 6% of the metro's land area) while 34%, 36%, and 22%, respectively, are contained in the region beyond 40km from this point (a region that contains 78% of the land area). Model-implied wages are on average 109 % higher in the center than in the periphery partly due to larger fundamental productivity

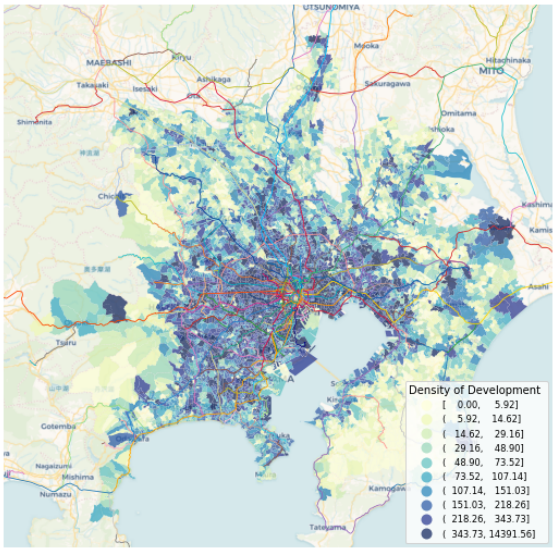
Table 5: Additional Parameters Used in Counterfactual with Road Congestion

Parameter	Value	Source
Elasticity of Travel Time w.r.t Traffic Flows (σ)	0.35	Allen and Arkolakis 2022
Route Heterogeneity Parameter (ρ)	50	Chosen
Fixed Travel Time ($\tilde{t}$)	6min	Estimated

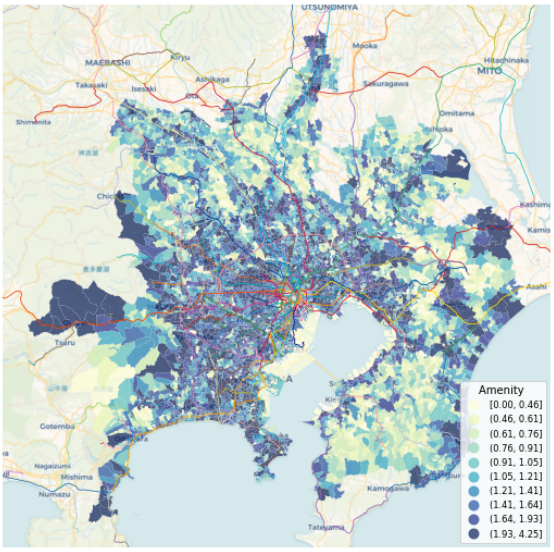
(80% higher), more capital deepening (60% higher floorspace) and agglomeration (contributing an additional 33% to the difference in productivity).



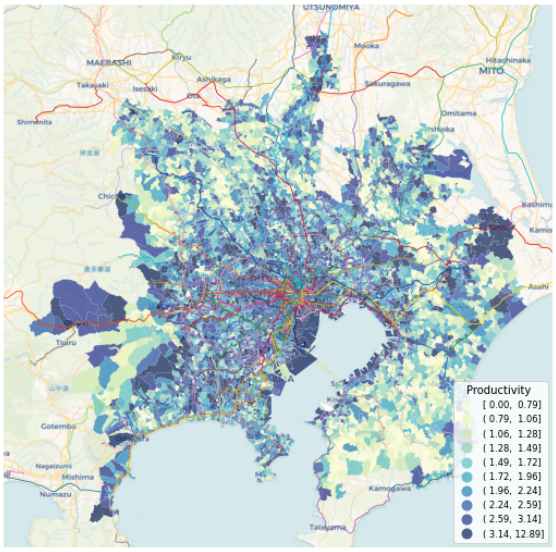
Floor Space



Density of Development



Amenities



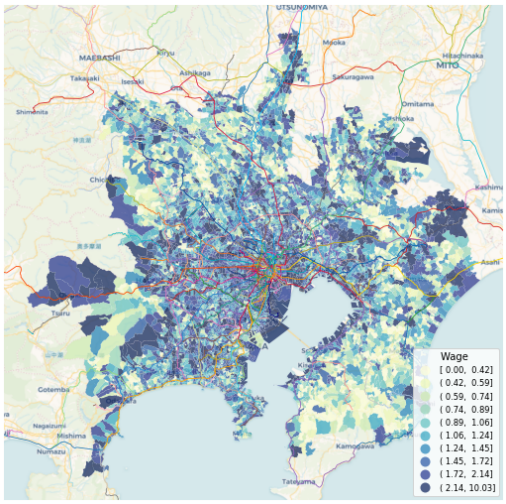
Productivity

Figure 4: Fundamentals and Floorspace Implied By 2020 Equlibrium

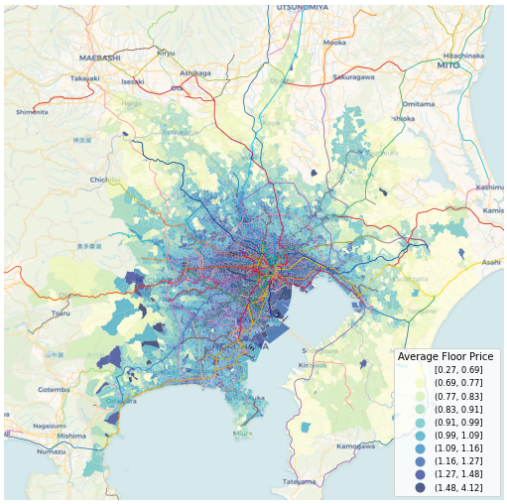
9.2 Removal of Tokyo's Train System

Using our calibrated model, we simulate Tokyo's economy without its train system.

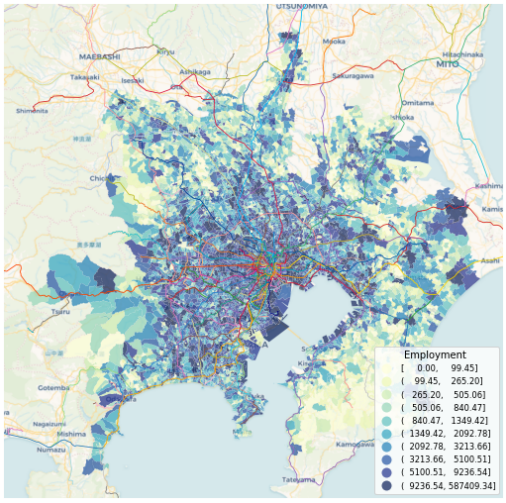
Specifically, we remove all choices involving train from agents' choice sets such that commuters can now only choose from walking, cycling, driving, or taking the bus. The new commuting cost



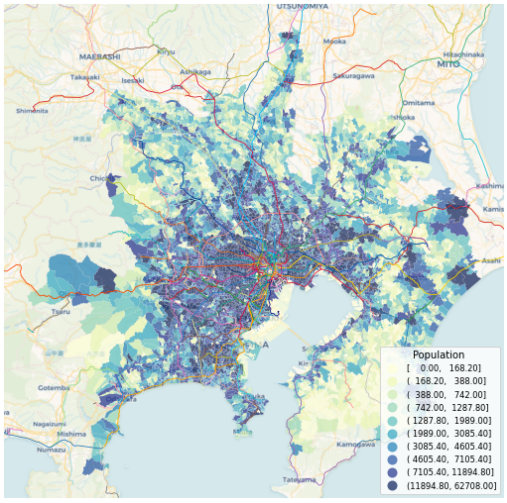
Wages



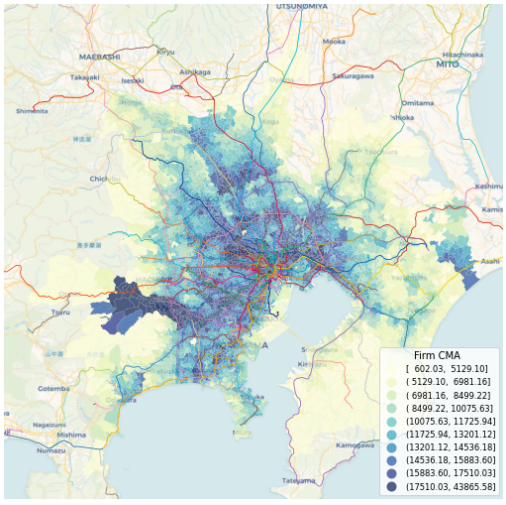
Floor Prices



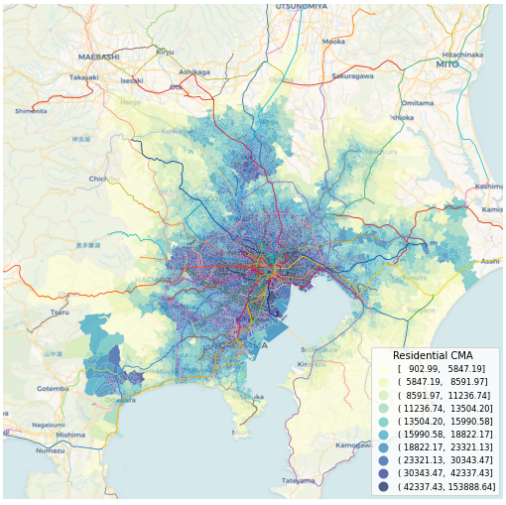
Employment



Residential Population

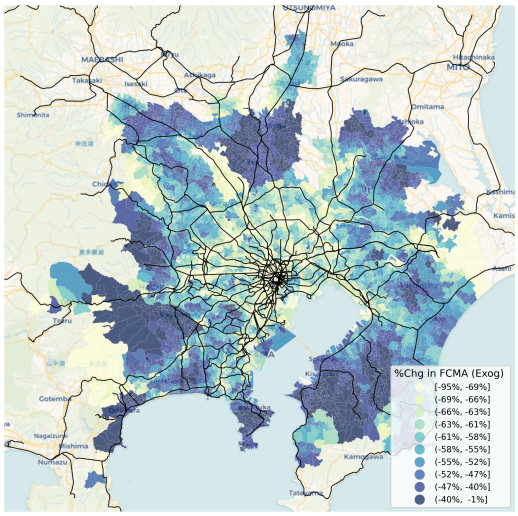


Firm CMA

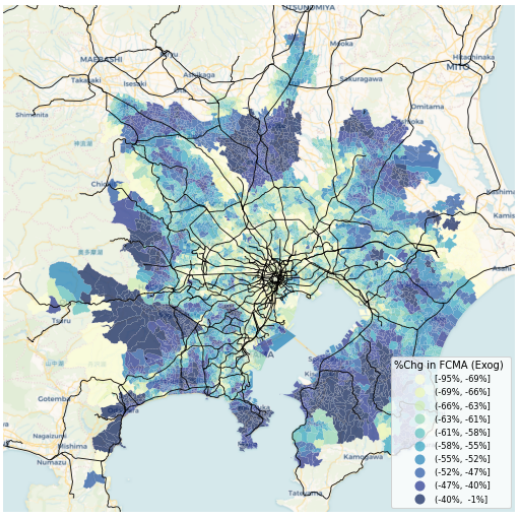


Residential CMA

Figure 5: 2020 Equilibrium



Exogenous Change to Firm CMA (% Change)



Exogenous Change to Residential CMA (% Change)

Figure 6: Shock to Commuting Access

between two locations i and j are then given by:

$$d_{ij} = \left[\sum_{m \backslash \text{train}} \exp(-\theta(\kappa_{0m} + \kappa_{1m} t_{ijm})) \right]^{-1/\theta} \tag{49}$$

where $\sum_{m\backslash\text{train}}$ denotes that the sum is taken over all modes except for the train.

We represent the shock to commuting access in Figure 6. These represent the change in commuting access if endogenous variables are held fixed at the 2020 equilibrium. When trains are removed, all regions lose commuting access with the median census tract declining by 58% in firm access and median censust tract declining by 61% in residential access. However, the decline in central Tokyo's firm access is 28% more than that of peripheries while the decline in residential access is 11% more than that of the peripheries. These discrepencies lead to shift in populalation and employment.

9.3 Counterfactual without Congestion

In the first counterfactual, we hold congestion (i.e. travel times), floorspace, and Tokyo's population fixed.

Given positive spillovers (μ_U and μ_A), there is no guarantee that the counterfactual will lead to a unique equilibrium. We use an equilibrium search rule such that we first solve for a counterfactual equilibrium without spillovers, and then use the results to initialize the search for an equilibrium with spillovers. The results of the counterfactual are shown in Figure 7. Removal of the train system leads to a shift in both employment and residential population to rural areas. Central Tokyo, which relies heavily on the train system, would lose 3% of its population, 12 % in employment, and 10% in land value while rural areas that are already car-dependent gain 8% in population, 20% in

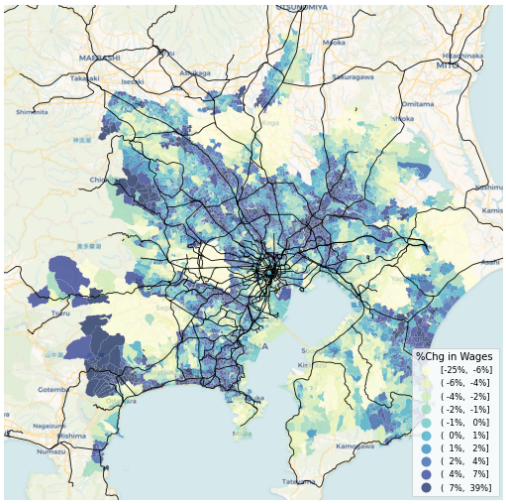
employment, and 4% in land values. Welfare declines by 33%.

9.4 Counterfactual with Road Congestion

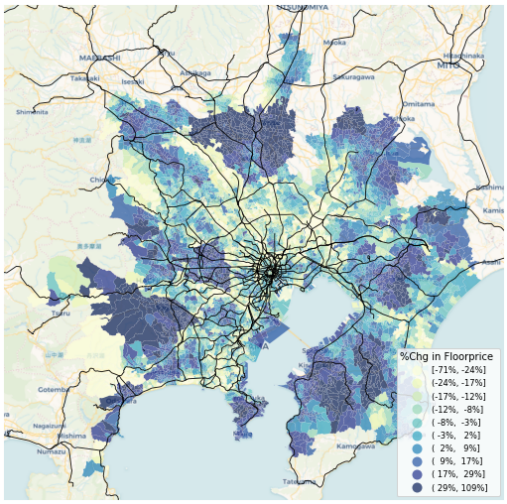
Work on simulations with road congestion is ongoing. Here, we show preliminary results using municipality-level data. Simulations that incorporate the impact of road congestion on car travel times imply that population and employment decline by 9% and 15% respectively in the city center while the peripheries gain 13% in population and 32% in employment. While increased congestion accounts for an additional 12% in automobile commute time on average (from 22 to 26 min), the spatial reorganization of the economy implies that commuters travel shorter distances and therefore automobile commute time decreases by 2 minutes on average (from 27.5 to 25.6 min) relative to the current 2020 equilibrium. Relative to the counterfactual without congestion, welfare declines by an additional 4 percentage points (37% in total). Our findings show quantitatively how automobile-focused urban areas lead to low-density "sprawl" as is observed in many American cities. However, there are several factors that will likely alter the results significantly. First, the current congestion model uses a parametrization that is likely a lower bound on the true impact of congestion. This is because the paramater σ, set to 0.35 (roughly halfway between between the range of estimates in Allen and Arkolakis 2022) was estimated as the elasticity of travel times with respect to traffic flows, whereas we parameterized 0.35 as the elasticity of the "time-in-traffic" (the endogenous component of travel time) with respect to traffic flows. Making the appropriate adjustment will likely increase the impact of congestion. Second, the current model does not allow for the endogenous adjustment of floorspace, which would likely lead to a more pronounced shift in population and employment. Third, the current model keeps the city's population fixed. Allowing the city's population to adjust (outward migration) has an immediate dampening effect on welfare, and also has a secondary welfare effect by reducing road congestion. Finally, the current model does not allow for the endogenous adjustment of road infrastructure, nor an accounting of the costs of construction. These factors will be incorporated in future work.

10 Conclusion

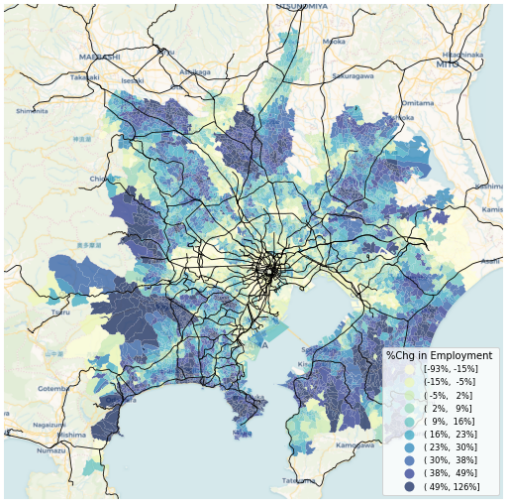
This paper has sought to evaluate the impact of Tokyo's train system on the distribution of economic activity and welfare. We estimated and calibrated an urban model of commuting using data on population, employment, land prices, travel times, commuter flows, and wages. We then performed a counterfactual simulation of Tokyo's economy without its train system and found substantial shifts away from urban centralization leading to a more dispersed population. The results suggest that the train system has a substantial impact on the distribution of economic activity and welfare in Tokyo. Future work will focus on incorporating road congestion into the model and endogenizing the density of development and using the expansion of the train network to estimate agglomeration paramters. We also seek to incorporate cost data in order to do a cost-benefit analysis.



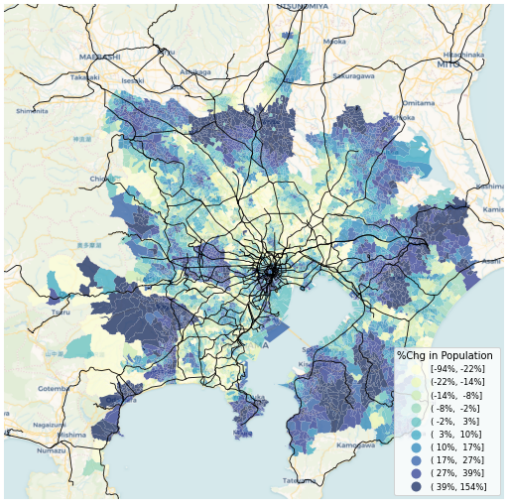
Changes in Wages



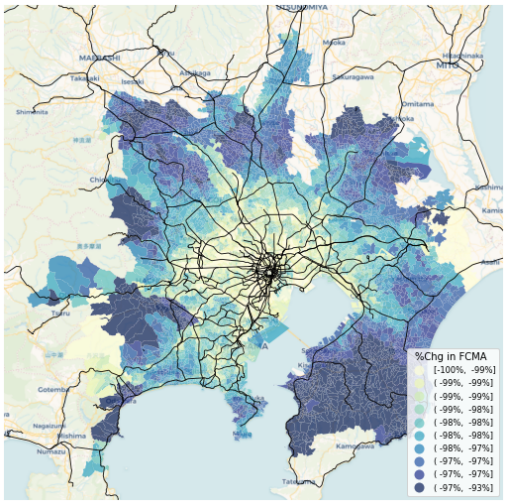
Changes in Floor Prices



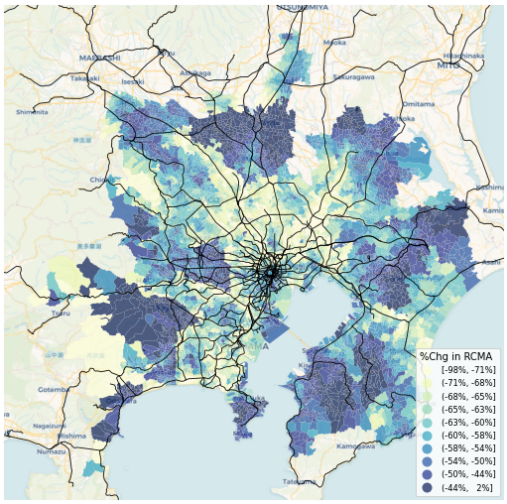
Changes in Employment



Changes in Residential Population



Changes in Firm CMA



Changes in Residential CMA

Figure 7: Changes in Variables from Train Removal with Road Congestion Fixed(% Change)

References

Ahlfeldt, Gabriel M. et al. (2015). "The Economics of Density: Evidence From the Berlin Wall". In: *Econometrica* 83.6. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA10876, pp. 2127–2189.

Allen, Treb and Costas Arkolakis (Aug. 1, 2014). "Trade and the Topography of the Spatial Economy". In: *The Quarterly Journal of Economics* 129.3. Publisher: Oxford Academic, pp. 1085–1140.

— (Nov. 1, 2022). "The Welfare Effects of Transportation Infrastructure Improvements". In: *The Review of Economic Studies* 89.6, pp. 2911–2957.

— (May 1, 2023). "Economic Activity across Space: A Supply and Demand Approach". In: *Journal of Economic Perspectives* 37.2, pp. 3–28.

Allen, Treb, Costas Arkolakis, and Xiangliang Li (Dec. 1, 2024). "On the Equilibrium Properties of Spatial Models". In: *American Economic Review: Insights* 6.4, pp. 472–489.

Asahi, Kenzo, Andrea Herrera, and Hugo E Silva (2023). "The effects of transport infrastructure on housing supply and the role of land-use regulation". In.

Banerjee, Abhijit, Esther Duflo, and Nancy Qian (June 1, 2020). "On the road: Access to transportation infrastructure and economic growth in China". In: *Journal of Development Economics* 145, p. 102442.

Baum-Snow, Nathaniel (2007). "Did Highways Cause Suburbanization?" In: *The Quarterly Journal of Economics* 122.2. Publisher: Oxford University Press, pp. 775–805.

Baum-Snow, Nathaniel, Loren Brandt, et al. (2017). "Roads, Railroads, and Decentralization of Chinese Cities". In: *The Review of Economics and Statistics* 99.3. Publisher: MIT Press, pp. 435–448.

Baum-Snow, Nathaniel, Matthew E. Kahn, and Richard Voith (2005). "Effects of Urban Rail Transit Expansions: Evidence from Sixteen Cities, 1970-2000 [with Comment]". In: *Brookings-Wharton Papers on Urban Affairs*. Publisher: Brookings Institution Press, pp. 147–206.

Berger, Thor and Kerstin Enflo (Mar. 1, 2017). "Locomotives of local growth: The short- and long-term impact of railroads in Sweden". In: *Journal of Urban Economics*. Urbanization in Developing Countries: Past and Present 98, pp. 124–138.

Bernard, Andrew B., Andreas Moxnes, and Yukiko U. Saito (Apr. 2019). "Production Networks, Geography, and Firm Performance". In: *Journal of Political Economy* 127.2, pp. 639–688.

Bernhofen, Daniel M. and John C. Brown (2005). "An Empirical Assessment of the Comparative Advantage Gains from Trade: Evidence from Japan". In: *The American Economic Review* 95.1, pp. 208–225.

Brooks, Leah and Zachary D. Liscow (Dec. 1, 2021). *Infrastructure Costs*. Rochester, NY.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (July 1, 2019). "The Costs of Agglomeration: House and Land Prices in French Cities". In: *The Review of Economic Studies* 86.4, pp. 1556–1589.

Correia, Sergio (Aug. 21, 2023). "REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects". In: *Statistical Software Components*. Publisher: Boston College Department of Economics.

Correia, Sergio, Paulo Guimarães, and Thomas Zylkin (Mar. 2020). "ppmlhdfe: Fast Poisson Estimation with High-Dimensional Fixed Effects". In: *The Stata Journal: Promoting communications on statistics and Stata* 20.1, pp. 95–115.

Davis, Morris A. and François Ortalo-Magné (Apr. 1, 2011). "Household expenditures, wages, rents". In: *Review of Economic Dynamics* 14.2, pp. 248–261.

Dingel, Jonathan I and Felix Tintelnot (2023). "Spatial Economics for Granular Settings". In.

Donaldson, Dave (Apr. 2018). "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure". In: *American Economic Review* 108.4, pp. 899–934.

Donaldson, Dave and Richard Hornbeck (May 1, 2016). "Railroads and American Economic Growth: A "Market Access" Approach *". In: *The Quarterly Journal of Economics* 131.2, pp. 799–858.

Duranton, Gilles, Geetika Nagpal, and Matthew Turner (May 2020). *Transportation Infrastructure in the US*.

Epple, Dennis, Brett Gordon, and Holger Sieg (June 2010). "A New Approach to Estimating the Production Function for Housing". In: *American Economic Review* 100.3, pp. 905–924.

Faber, Benjamin (July 1, 2014). "Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System". In: *The Review of Economic Studies* 81.3, pp. 1046–1070.

Fajgelbaum, Pablo D., Cecile Gaubert, et al. (July 2023). *Political Preferences and the Spatial Distribution of Infrastructure: Evidence from California's High-Speed Rail*.

Fajgelbaum, Pablo D. and Edouard Schaal (2020). "Optimal Transport Networks in Spatial Equilibrium". In: *Econometrica* 88.4. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15213, pp. 1411–1452.

Free, Dan (Nov. 7, 2017). *Early Japanese Railways 1853-1914: Engineering Triumphs That Transformed Meiji-era Japan*. Reprint edition. Tuttle Publishing. 288 pp.

Glaeser, Edward, Matthew E. Kahn, and Jordan Rappaport (2008). "Why Do The Poor Live In Cities? The Role of Public Transportation". In: *Journal of Urban Economics*. Accepted: 2009-05-14T18:30:49Z Publisher: Elsevier.

Glaeser, Edward L. and Matthew E. Kahn (Feb. 2001). *Decentralized Employment and the Transformation of the American City*.

— (Jan. 1, 2004). "Chapter 56 - Sprawl and Urban Growth". In: *Handbook of Regional and Urban Economics*. Ed. by J. Vernon Henderson and Jacques-François Thisse. Vol. 4. Cities and Geography. Elsevier, pp. 2481–2527.

Gonzalez-Navarro, Marco and Matthew A. Turner (Nov. 1, 2018). "Subways and urban growth: Evidence from earth". In: *Journal of Urban Economics* 108, pp. 85–106.

Hirooka, Haruya (2000). "The Development of Tokyo's Rail Network". In.

Kazunobu, HAYAKAWA et al. (2021). "High-speed Rail and the Spatial Distribution of Economic Activity: Evidence from Japan's Shinkansen". In.

LaPoint, Cameron (2020). "You Only Lend Twice: Corporate Borrowing and Land Values in Real Estate Cycles". In: *SSRN Electronic Journal*.

Meyer, John R. et al. (Jan. 1, 1959). *Economics of Competition in the Transportation Industries*. Cambridge Massachusetts London, England: Harvard University Press. 377 pp.

Meyer, John Robert (Jan. 1, 1965). *The Urban Transportation Problem*. Pencil Margin Notes edition. Cambridge, Mass: Harvard Univ Pr. 448 pp.

Miyauchi, Yuhei, Kentaro Nakajima, and Stephen Redding (Feb. 2021). *The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data*. w28497. Cambridge, MA: National Bureau of Economic Research, w28497.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg (Dec. 2018). "Commuting, Migration, and Local Employment Elasticities". In: *American Economic Review* 108.12, pp. 3855–3890.

Redding, Stephen J. and Daniel M. Sturm (Dec. 2008). "The Costs of Remoteness: Evidence from German Division and Reunification". In: *American Economic Review* 98.5, pp. 1766–1797.

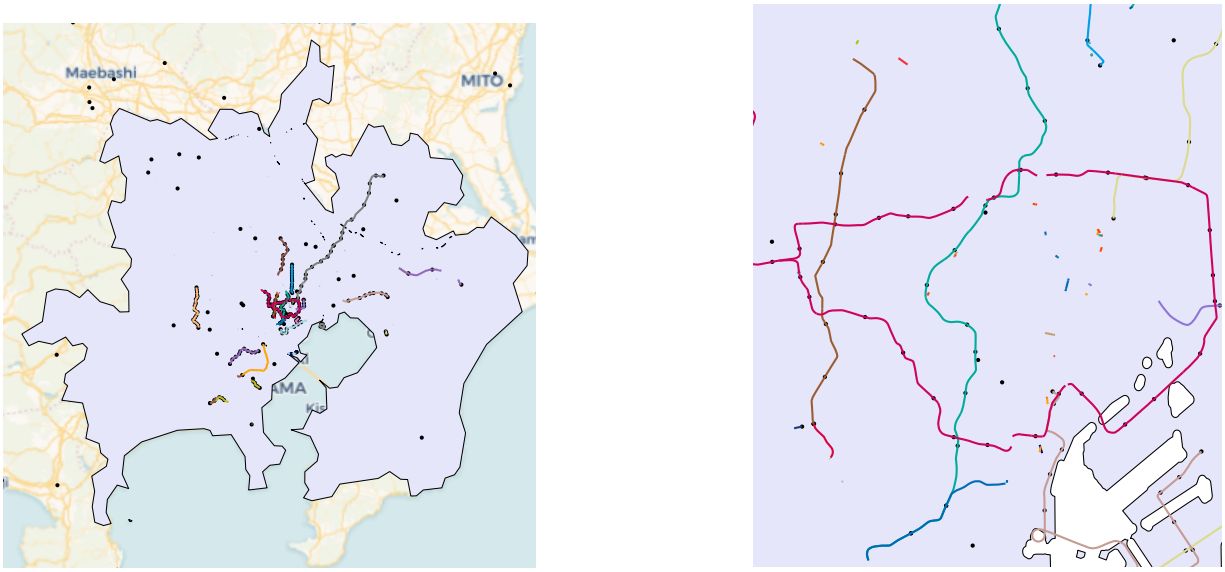
Tsivanidis, Nick (2022). "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio". In.

Tsivanidis, Nick (2023). "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio". In.

Valentinyi, Akos and Berthold Herrendorf (2008). "Measuring Factor Income Shares at the Sector Level". In: *Review of Economic Dynamics* 11.4. Publisher: Elsevier for the Society for Economic Dynamics, pp. 820–835.

A Estimating Spillover Parameters

Between 1990 and 2020, there were 75 train-line modifications that added approximately 500 km of new track to the system, amounting to a ∼10 percent increase in track length. These changes, which can be seen in Figure 8, spanned both the central Tokyo area as well as the peripheries. In future work, we will use this variation to estimate spillover parameters and also to assess the fit of the model to the data.



Greater Tokyo Area Central Tokyo

Figure 8: Changes to Train Lines from 1990-2020

Colored lines are train lines while black dots are stations. Shaded region is the Greater Tokyo Area. Data is from the Ministry of Land, Infrastructure, Transport, and Tourism (MLIT)

Given our expressions for productivities and amenities as described in the previous section, the next step is to estimate spillover parameters using an instrumental variables strategy. After we have obtained productivities using equation (43), our expression for spillovers (equation (17)) then implies the following regression equation:

$$\ln A_j = \ln \bar{A}_j + \mu_A \ln L_{Fj} \tag{50}$$

$$\ln U_i = \ln \bar{U}_i + \mu_U \ln L_{Ri} \tag{51}$$

In the above equations, the exogenous productivities and amenities, $\ln \bar{A}_j$ and $\ln \bar{U}_i$, are residuals, and therefore there are two problems that need to be addressed before proceeding to the regression. First, there is no guarantee that the residuals are mean zero. This problem can be addressed by adding a constant to the right hand side of the above equations. The second issue is that

our model implies that the residuals are correlated with the regressors. Higher exogenous amenities lead to a higher supply of residents, and likewise, higher exogenous productivies lead to a higher supply of workers. To address this issue, we can use instruments for the regressors. Ahlfeldt et al. 2015 compute the change in exogenous productivities and amenities between the pre-war period and the division of Berlin to eliminate location-specific fixed effects to amenities and productivities. They then use indicators for distance grid cells from the pre-war CBD in Berlin as instruments, arguing that the systematic change in the gradient of economic activity in West Berlin relative to the pre-war CBD following division is explained by the mechanisms of the model rather than by systematic changes in the pattern of structural residuals. The instruments are also relevant because locations closer to the pre-war CBD experienced a drop in employment and population relative to locations further away.[29] Using a similar approach, Tsivanidis 2023 use instruments for the change in RCMA and FCMA when population and employment are fixed at their initial (i.e. predetermined) values.

In the same spirit as Tsivanidis 2023, we will use the change in RCMA and FCMA when variables other than the travel times are held fixed to the pre-expansion values. Specifically, our instrument for estimating residential externalities will be

$$\tilde{\Delta} \ln \Phi_{Ri} \equiv \ln \sum_j (w_{j,pre}/d_{ij,post})^\theta - \ln \sum_j (w_{j,pre}/d_{ij,pre})^\theta$$

where the subscript *pre* denotes that the variable is from the pre-expansion period while the subscript *post* denotes that the variable is from the post-expansion period. Similarly, the instrument for estimating production externalities is given by:

$$\tilde{\Delta} \ln \Phi_{Fj} \equiv \ln \sum_i (U_{i,pre} r_{Ri,pre}^{\beta-1}/d_{ij,post})^\theta - \ln \sum_i (U_{i,pre} r_{Ri,pre}^{\beta-1}/d_{ij,pre})^\theta$$

where we have once again used predetermined variables for the potentially endogenous variables, $U_{i,pre}$ and $r_{Ri,pre}$, to construct the instruments.

In short, we assume that changes in travel times as uncorrelated with changes in the structural residuals, and therefore interact these changes with pre-determined variables to constuct variables that are both uncorrelated with the residuals and have relevance with respect to the explanatory variable (through improved market access).

B Travel Time Calculation

Travel times for each mode are computed using GIS and network analysis software between each census tract. We then calculate a factor for each pair of census tracts such that the average travel times computed across origin and destination census-tract pairs within commuter-zone pairs are

[29] Ahlfeldt et al. 2015 also estimate parameters that govern the spatial scale of spillovers. Their estimation strategy therefore requires more instruments. While we may seek to estimate the spatial scale of spillovers in future work, for now we will proceed with the simpler version.

equated to the average travel times reported in the survey.[30] For census-tract pairs where we cannot compute this factor (because there are no reported trips), we use the predicted times from a regression of reported time on computed times. In this section, we describe the details of these steps.

B.1 Computing Travel Times for Cars, Walking, Cycling, and Buses

We use Openroute Service to compute travel times between the centroids of every pair of regions in Tokyo (for determining the automobile travel time) and also between each centroid and nearby train stations (for the last-mile mode of transport). Openroute Service is an opensource tool that utilizes the street network provided by OpenStreetMap to calculate routes between locations.[31] The tool typically assumes that vehicles travel at the speed limit of a road segment.[32] It is thus generally more accurate for travel times in low-traffic settings when the speed limit is a good proxy for the actual speed of travel. Whenever the centroid of a location does not land on a road segment, the tool will use the closest road segment within a 350-meter radius of the centroid.[33]

B.2 Computing Train Travel Times

We use longitudinal data on the train network provided by the MLIT to compute travel times. The data includes geographical information on each train line and station of each train line. We use this geographical data to construct a graph on Python's NetworkX library. Each station of a train line is a node and each segment of the train line between two neighboring stations forms an edge. We assume that passengers can transfer between any two stations within 1 km of each other by walking. If they do so, they must wait 4 minutes on the next train platform in addition to the time it takes to walk, assuming a walking speed of 4 km/hr using a straight-line distance.

If a station is a hub for multiple train lines, each train line receives its own node, and these nodes are connected by edges that represent the time it takes to transfer between train lines, following the same rules as stated above.

Each edge in the graph is weighted by the time it takes to travel between the two nodes. This time is calculated by dividing the distance (provided by the GIS map) between the two nodes by the effective train speed, which were looked up for 24 of the major train lines and were assumed to be 40 km/hr for other train lines.[34]

Households are located at the centroid of each district, and are able to travel to their nearest stations by either walking or taking the bus. If they walk, they once again travel at 4 km/hr using the straight-line distance, whereas if they travel by bus, they will travel through the street network

[30] This approach follows Miyauchi, Nakajima, and S. Redding 2021.

[31] OpenStreetmap is a free and open geographic database of maps (including street maps) from across the world

[32] The documentation of Openroute Service writes "The travel time is calculated for each segment by using speed-limits for different waytypes and adjusting them for different surfaces or surface qualities of the road."

[33] There were two tracts however where the centroids were not within 350 meters of a road segment. For these tracts, we manually selected a point on a nearby road segment instead and used this point instead.

[34] 40 km/hr was chosen because this was the typical speed that we observed for the major train lines

at half the speed of a car in addition to paying a fixed time cost of 6 minutes, which represents the additional time from walking to the bus stop and waiting for the bus.

B.3 Incorporating Reported Travel Times from Survey

As discussed in the main text, we believe times reported in the travel survey (*reported time* henceforth) are more accurate than the computed times (from Openroute Service or NetworkX) because the reported times do not rely on ad-hoc assumptions about travel speeds used in the computed time, and they should also take into account congestion, lastmile travel times, and other nuances of travel that may not be captured by the computed time. The two drawbacks of the reported time are that travel times are reported only for pairs of locations where trips have been reported and the reported time is only reported at the geographic unit of the commuter-flow data (about 15 thousands residents, referred to as the "commuting district" in the main body and "small-zone" in data manual) as opposed to the census-tract level (about 2 thousand residents) which is the level of the model analysis. [35]

In order to retain the accuracy of reported travel time while having a full matrix of travel times between census tracts for each mode, we take the folloiwng steps to incorporate both the reported and computed times. First, we calculate the average computed travel time between small zones by taking the simple average of computed travel times across pairs of census tracts between the any two small zones. When the data is available (i.e. there are reported trips at the small-zone level), we calculate the ratio between the computed travel time and the reported travel time. We then use this factor to scale the computed time at the small-zone level. Since the reported time is missing for many pairs of small zones, we then aggregate the computed and reported times to three levels of geographical aggregation. At the highest level of aggregation, each region contains on average 250 thousand residents (less than one percent of the population). We then compute the ratio between the computed travel time and the reported travel time for each level of aggregation and use this factor to scale the computed time at the census-tract level, using the lowest level of aggregation that is available.[36] From an initial dataset with 1.6 percent of pairs filled in (these are the diade-mode combinations that have at least one trip), the above procedure fills in an additional 42 percent of location-pair and mode combinations. The remaining observations are then filled in by running a regression of reported time on computed time for each mode and using the predictions from this regression model. These regressions and corresponding scatter plots are shown in Table 6 and Figure 9. For the predictive regressions, we use a simple linear model with a constant, because aside from an intercept (reflecting a fixed cost of travel), our prior was that reported and computed times should scale linearly. Scatter plots from 9 support this hypothesis.

[35] Another reason that the computed time is typically needed is that it is used to compute travel times for counterfactual transportation networks, though in the present analysis, we remove the entire train network, and so an alternative train network did not need to be computed.

[36] Thus, these computed times are scaled such that the weighted-average travel time at the higher-level of aggregation matches that of the reported time, whenever the time has been reported.

C Details on Data and Assumptions

In this section, we describe additional details about the data and what assumptions are made to transform the data from the raw data to the processed data that is used in the analysis.

C.1 Merging Employment and Population Censuses

Both the employment and population census data appear at the block levels, but they do not utilize the same block identification codes. Because both datasets contain the address of the block, we merge the two datasets by geocoding the addresses using geocode software, and then merging the blocks using a spatial merge with GIS tools.[37] We geocode the addresses using the *CSV Address Matching Service*, a tool from Tokyo University's Center for Spatial Information Science.

The same challenge arises with the 1995 population data and 1996 employment data because the definitions of blocks and census tracts are not consistent across time. We therefore rely once again on the *CSV Address Matching Service* by looking up the coordinates from the addresses that are listed on the website[38]

Specifically, we use the following procedure. After looking up the coordinates of the 1995 population tracts, 1996 employment tracts, and 2021 employment tracts, we proceed by merging the 2020 population tracts with the 2021 employment coordinates, then merge the resulting districts with the 1996 employment coordinates, and then merge the resulting districts with the 1995 population coordinates. When we merge the 2020 population tracts (which are polygons in GIS jargon) with the 2021 employment coordinates (points in GIS jargon), we first perform a spatial inner join between the polygons and points. Thus, all coordinates that are contained in a 2020 population tract are joined. Then polygons that did not match with any points are matched with the closest employment point. If the nearest employment point has already been matched to another polygon, then the two polygons are merged (i.e. aggregated) together. If there are employment points that do not join with residential tracts, they are dropped. We then repeat this procedure to merge these new polygons with the 1996 employment cooordinates, and then again with the 1995 population coordinates.

We continue to use the term "tract" for these new district definitions, because the difference the definitions are negligible.

[37] Directly merging the datasets using the addresses is avoided because the addresses are not standardized across the two datasets.

[38] These addresses are also not consistent over time. However, the *CSV Address Matching Service* is able to geolocate addresses that used earlier address names. For each address lookup the service provides an estimate of the accuracy: whether the coordinate is correct up to the prefecture, municipality ("shichooson" + "tokubetsuku"), or census tract ("ooaza"). We found that 0.2 percent of the 2021 employment addresses were matched up to the municipality level, 0.2 percent of the 1996 employment addresses were matched to this level, and 5.8 percent of the 1995 population addresses were matched to this level. All other addresses were matched to the census-tract level. Our census

	Train	Weighted	Weighted (capped)	Car	Weighted	Weighted (capped)
Computed Time	0.537	0.564	0.573	1.551	1.638	1.742
	(0.002)	(0.002)	(0.002)	(0.011)	(0.012)	(0.010)
Intercept	26.792	23.062	22.519	13.375	8.760	7.730
	(0.154)	(0.138)	(0.141)	(0.184)	(0.148)	(0.126)
N	115648	115648	113946	49320	49320	48096
R-squared	0.47	0.54	0.52	0.52	0.60	0.57

	Bus	Weighted	Weighted (capped)	Cycle	Weighted	Weighted (capped)
Computed Time	1.398	1.427	2.430	0.672	0.635	0.933
	(0.062)	(0.091)	(0.065)	(0.039)	(0.048)	(0.014)
Intercept	29.695	27.379	20.618	13.743	11.290	8.643
	(0.521)	(0.655)	(0.460)	(0.580)	(0.440)	(0.139)
N	7952	7952	7814	14617	14617	14107
R-squared	0.21	0.22	0.25	0.23	0.26	0.30

	Walk	Weighted	Weighted (capped)
Computed Time	0.141	0.201	0.264
	(0.017)	(0.011)	(0.008)
Intercept	19.539	12.100	10.899
	(0.607)	(0.218)	(0.157)
N	9441	9441	8783
R-squared	0.09	0.21	0.24

Table 6: Regressions of Reported Time on Computed Time

OLS regressions of reported time on computed time for each mode. The unit of observation is the "small" zone, defined by the survey to contain on average 15 thousand residents. Weighted regression use frequency weights based on the number of trips of that pair. "Weighted (capped)" restricts the sample to pairs of locations that are within the 99th percentile of the distribution of computed travel times by mode. Data on commuter trips and reported times are obtained by Tokyo's Person-Trip Survey. Computed travel times by car, bicycle, and walking are calculated using OpenRoute Services while travel times by train are calculated using the train network and OpenRoute Service. Computed bus travel time is the same as the car travel time. See Data section and Appendix for details. Robust standard errors in parentheses.

C.2 Gravity Estimation

The gravity model is estimated using trips reported on the Tokyo Person Trip Survey data. For the righthand-side variable (travel time), we use the methodology discussed in Section B, particularly Section B.3, where we fill in missing times using a prediction using the GIS-computed times.

The baseline Poisson Pseudo-Maximum-Likelihood (PPML) specification uses the commuter flows from the survey on the left-hand-side and travel times calculated from the procedure above. Observations whose travel times were calculated using the prediction algorithm are weighted proportionally to the R^2 of the prediction regression (while other observations have a weighting of 1). To check the robustness of our findings, we also estimate all combinations of specifications along the following dimensions. First, we vary our right-hand-side variable to incorporate the computed time, the reported time, and the "hybrid" time, which are the times from our baseline.

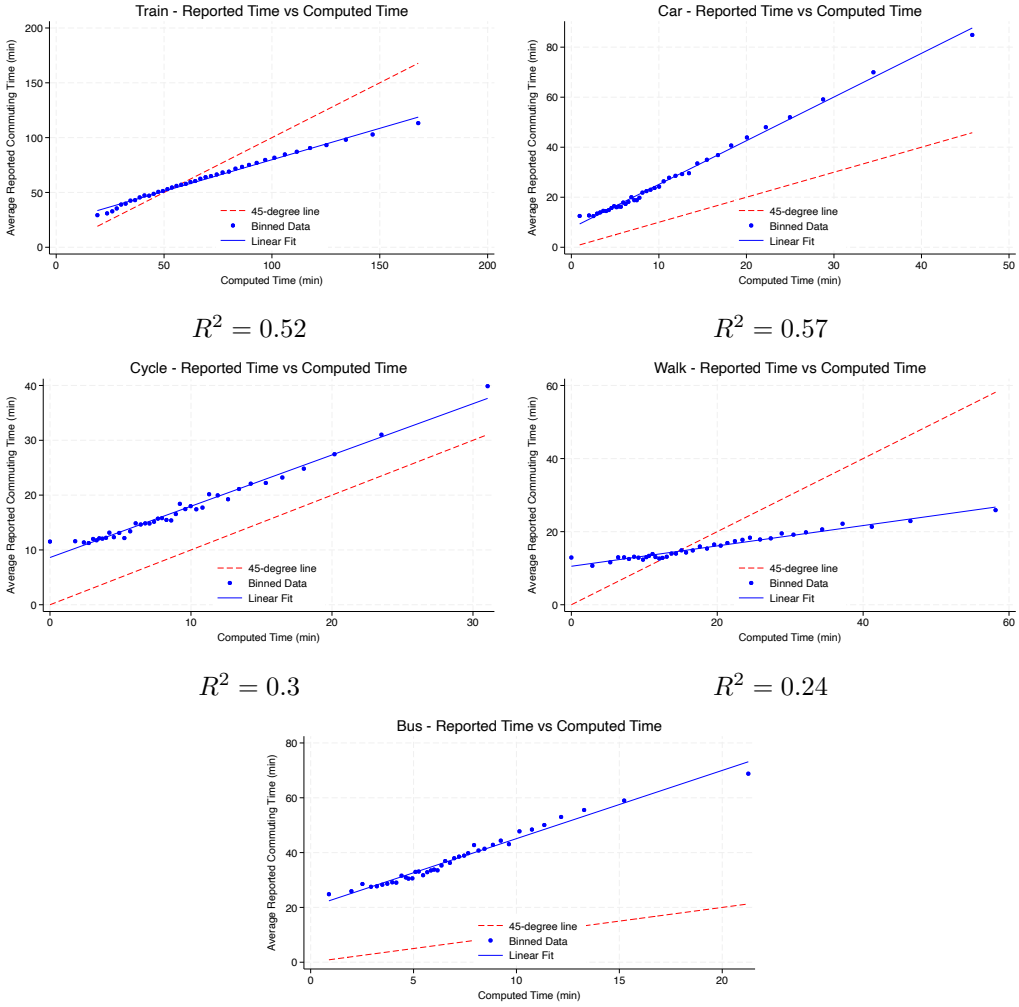


Figure 9: Binscatter and OLS Fit of Reported Time on Computed Time

Binscatter plots and OLS fit from a regression of reported time on computed time for each mode. The R^2 of the regression is shown in the caption. Data on commuter trips and reported times are obtained by Tokyo's Person-Trip Survey. The unit of observation is the "small" zone, defined by the survey to contain on average 15 thousand residents. Computed travel times by car, bicycle, and walking are calculated using OpenRoute Services while travel times by train are calculated using the train network and OpenRoute Service. Computed bus travel time is the same as the car travel time. The regression is weighted by the number of trips. We restrict the regression and binscatter to the pairs of locations that are within the 99th percentile of the distribution of computed travel times by mode. See Data section and Appendix for details.

C.2.1 Cleaning Data for Gravity Estimation

Prior to estimating gravity, we take the following steps to clean the data:

- We found 65 origin-destination observations in the bilateral flows data where the survey reported that commuters walked when the reported commuted times are unrealistically short. We believe these trips were misported as walking trips when they were in fact trips taken by train.[39] For these observations, we recoded the mode as train trips.

C.3 Merging Commuter Flows Data with Census Data

C.4 Model Fit

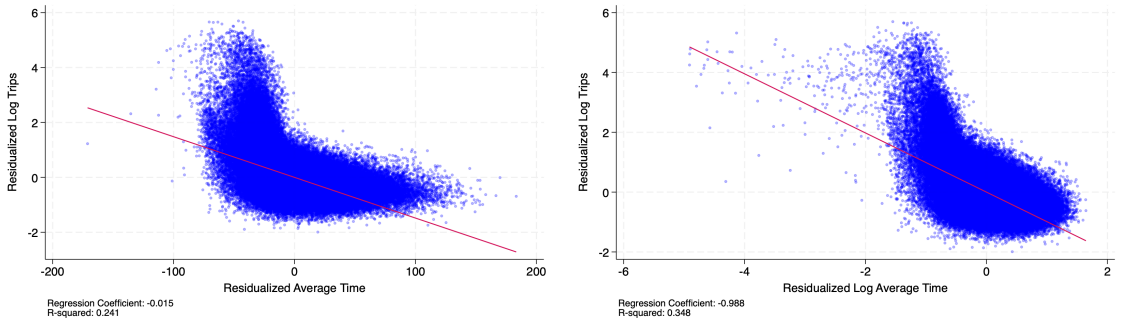
We perform a number of validation checks to assess the model's fit. First, we compare the model-implied wages to the municipality-level taxable earnings that are published by each municipality. We find a correlation of 0.71 suggesting the model has explanatory power for municipality wages.

Second, we assess the model fit of the exponential commuting-cost assumption by analyzing the residualized scatter plot from the commuter gravity equation. The scatter plot after partialling out origin fixed effects and destination fixed effects from the dependent variable (log commuter flows) and the independent variable (average commuting time) can be found in Figure 10. Next to it, we also present the scatter plot with a log-log specification. Both plots, as expected, show large negative relationships between commuting trips and commuting times with substantial explanatory power, with over half of the variation in residualized log trips being explained by the residualized log travel times. While the semi-log specification implied by our model has substantial explanatory power (R^2 of 0.6), the log-log specification not only has a higher coefficient of determination (R^2 of 0.7), but also appears to exhibit more uniformity in fit across the distribution (see figure). This observation suggests it may be appropriate to adopt a power-function specification for commuting cost. This is a worthwhile area for future iterations of this paper.

In future work, we also plan to perform additional validation checks. For example, Ahlfeldt et al. 2015 compare the cumulative distribution function of commuters across travel time bins (e.g. 0-5 minutes, 5-10 minutes and so on) in the model and commuter-flow survey data from a number of years. In future work, we plan to perform this exercise in our setting.

In addition, we plan to compare the model-implied population, employment, wage, and land-prices changes between 1995-2020 with the actual changes that occured in the 1995-2020 period. This exercise is the strongest test of our model, as it allows us to assess the model's ability to predict observable variables when the transporation network changes.

[39] It is expected that many respondents misreported train trips as walking trips because reporting a train trip takes significantly more time to report. A respondent reporting a train trip is required to report every transfer station and train line used along the route, which can be time consuming, relative to simply reporting that one walked.



Log Commuting Trip vs Average Commuting Time

Log Commuting Trip vs Average Commuting Time

Figure 10: Residualized scatter plots from commuter flow regressions

Both the independent and dependent variables are first regressed on origin fixed effects and destination fixed effects. The $R-squared$ denotes the coefficient of determination from the residualized regression and therefore represents the explanatory power from travel time, indepndent of origin fixed effects and desintination fixed effects. Origin-destination combinations with no trips are not included because the natural log of zero is undefined. Data on commuter trips are obtained by Tokyo's Person-Trip Survey of approximately 630,000 individuals. Travel times are a weighted average travel times between cars and public transport where the weights are the share of trips taken by each mode in a particular municipality. Travel times by car are calculated using OpenRoute Services while travel times by public transport are calculated using NetworkX in addition to the geospatial information on the train network. See Data section and Appendix for details.

C.5 Additional Information on the Official Land Price Announcements

A land plot in the data is observed for an average of 17 years (median 15 years). Given that there exists turnover in the sample of surveyed land plots, there is a question as to whether a balanced or unbalanced sample should be used for the analysis. On the one hand, a balanced panel would allow for a more reliable measure of the *change* in land prices, which would be important for the estimation of the spillover parameters as we exploit variation across time(see section on estimation in section 7). On the other hand, given that the cross-sectional land plots are selected to be representative of a municipality, the unbalanced panel are likely better for studying variation in land prices across regions, which is an area of interest for the counterfactuals. We therefore plan to use the balanced panel for estimating the spillover parameters while using the unbalanced panel for the counterfactuals in future work. For the preliminary draft, we use the balanced panel for all analysis. Regardless of the methodology, this concern is minor given that the average tenure of a land plot is 17 years, and thus turnover is low. In the baseline analysis, we drop land plots that are not observed in all years and aggregate the land prices to the census tract level. For robustness, we repeat the analysis by aggregating without dropping land plots that fall out of the panel.

D Data Terminology and Translations

In this section, we describe the terminology used in the data section, so that the reader can understand where they can find the data used in the paper.

1. Population Census: Kokusei-chosa
2. Block level refers to cho-cho-moku-level in Japanese.
3. Census tract level refers to the Owaza-level in Japanese.
4. The Economic Census, which provides employment figures, refers to Keizai-toukei in Japanese.
5. The Official Land Price Announcements, which are used for land prices, are called Chika-koji in Japanese.