

NBER WORKING PAPER SERIES

THE ECONOMICS OF SPATIAL MOBILITY:
THEORY AND EVIDENCE USING SMARTPHONE DATA

Yuhei Miyauchi
Kentaro Nakajima
Stephen J. Redding

Working Paper 28497
<http://www.nber.org/papers/w28497>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2021, Revised July 2025

We would like to thank the editor, the anonymous referees, Gabriel Ahlfeldt, Milena Almagro, Daniel Sturm, Gabriel Kreindler, Tobias Salz and conference and seminar participants for helpful comments. We are grateful to Takeshi Fukasawa, Peter Deffebach, Atsushi Yamagishi, Bingqing Yang, Yun-Ting Yeh, and Chenming Zhang for their excellent research assistance. The usual disclaimer applies. “Konzatsu-Tokei (R)” Data refers to people flow data constructed from individual location information sent from mobile phones under users’ consent, through applications provided by NTT DOCOMO, INC (including mapping application Docomo Chizu NAVI and MyDaiz). Those data are processed collectively and statistically in order to conceal private information. Original location data is GPS data (latitude, longitude) sent every five minutes (minimum), and it does not include information to specify individual. The copyrights of all tables and figures presented in this document belong to ZENRIN DataCom CO., LTD. We also acknowledge Yaichi Aoshima at Hitotsubashi University for coordinating the project with ZENRIN DataCom Co., LTD.; Murata Foundation, Heiwa Nakajima Foundation, The Kajima Foundation, Obayashi Foundation, JSPS KAKENHI (Grant No. 21H00703), RIETI, and the Hitotsubashi University for their financial support; CSIS at the University of Tokyo for the joint research support (Project No. 954). We acknowledge Sanko Estate Company Ltd. for providing us with the office market data. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Yuhei Miyauchi, Kentaro Nakajima, and Stephen J. Redding. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data
Yuhei Miyauchi, Kentaro Nakajima, and Stephen J. Redding
NBER Working Paper No. 28497
February 2021, Revised July 2025
JEL No. R2, R3, R41

ABSTRACT

We develop a tractable quantitative framework for modelling the rich patterns of spatial mobility observed in smartphone data. We show that travel is frequently undertaken as part of a travel itinerary, defined as a journey starting and ending at home that can include more than one intermediate stop on a given day. We show that these travel itineraries provide microfoundations for consumption externalities and generate both complementarity and substitutability between locations. We show that the consumption externalities implied by travel itineraries are central to matching quasi-experimental evidence from the shift to WFH. We find that these consumption externalities are key drivers of the agglomeration of economic activity in central cities and shape the relative welfare gains from alternative transport improvements in favor of investments in central cities.

Yuhei Miyauchi
Boston University
miyauchi@bu.edu

Kentaro Nakajima
Hitotsubashi University
nakajima.kentaro@gmail.com

Stephen J. Redding
Stanford University
Department of Economics
and CEPR
and also NBER
reddings@stanford.edu

1 Introduction

Smartphone data have the potential to revolutionize our understanding of spatial mobility, because they measure travel with high spatial and temporal resolution. We use these data to show that travel frequently occurs as part of a *travel itinerary* – a journey starting and ending at home that can include more than one intermediate stop on a given day. Travel itineraries allow agents to visit multiple destinations at a lower total travel cost than implied by separate bilateral journeys to each destination.¹ Travel itineraries matter economically, because they create consumption externalities, and give rise to both complementarity and substitutability between locations. As one location becomes a more attractive destination, and experiences increased foot traffic, this makes visits to other nearby locations that are convenient intermediate stops more attractive (complements). At the same time, this redistribution of foot traffic makes visits to other locations that are further away and inconvenient intermediate stops less attractive (substitutes). These consumption externalities shape the organization of economic activity within cities, because they provide a force for agglomeration. When varieties of non-traded services are clustered closer together, or nearby home and work, they can be consumed at lower total travel costs. These consumption externalities also influence the impact of public policy interventions, because an intervention in one location is transmitted to other locations through the network of travel itineraries.

Although the empirical relevance of travel itineraries is known, developing tractable quantitative models of them has proved challenging. First, from an empirical perspective, most population censuses report bilateral commuting travel, but do not measure other types of travel. Although travel diary surveys report other types of travel, they are typically small-sample surveys rather than censuses, can be collected relatively infrequently, and do not always report travel sequences with fine spatial resolution across multiple days. Second, from a theoretical perspective, travel itineraries involve choosing ordered sequences of locations. With \mathcal{N} locations, there are $2^{\mathcal{N}}$ possible combinations of locations to visit, and many more sequences in which to visit them. In empirical applications with hundreds of locations or more, it is computationally infeasible to compute value functions for every sequence. One of our main contributions is to develop a tractable quantitative framework for modelling travel itineraries, which overcomes this challenge of a high-dimensional state space using importance sampling. We show that our quantitative framework is empirically successful in capturing the rich patterns of spatial mobility observed in our smartphone data.

We implement our quantitative framework using smartphone data for the Greater Tokyo

¹While we focus on personal travel, similar travel itineraries arise in other contexts in economics, including for example choices of intermediate stops for trucks, public transit and ships. A distinctive feature of personal travel itineraries is that they typically start and end at home on a given day.

metropolitan area. Our data come from a major smartphone mapping application in Japan (*Docomo Chizu NAVI*, *MyDaiz*), which records the Geographical Positioning System (GPS) location of each device every 5 minutes (at the highest frequency). We have anonymized data from 2015, before the COVID-19 pandemic, and data through the middle of 2024 afterwards. We measure each location visited by a user using a “stay,” which corresponds to no movement within 100 meters for 15 minutes. We categorize stays as home, work, and non-work. We define “travel itineraries” as a sequence of stays that starts and ends at home on a given day. A substantial fraction of these travel itineraries involve multiple intermediate stays (“trip chains”), such as stopping at a bar and a restaurant on the way home after work, or visiting multiple stores as part of a shopping expedition during the weekend.

We use our smartphone data to provide quasi-experimental evidence of consumption externalities. First, we examine the shift to working from home (WFH) in the aftermath of the COVID-19 pandemic. Consistent with evidence from other countries, we find a decline in work trips into Central Tokyo following the shift to WFH. Additionally, we find that this decline in work trips goes hand-in-hand with a reduction in non-work trips, which is consistent with consumption externalities from work travel, and with the idea that these two types of travel are jointly determined. Second, we report event-study estimates of the impact of the closure of large retail stores (floor space of more than 5,000 meters squared). In the months immediately afterwards, we find a decline in non-work stays in both the 250-meter grid cell containing the closed store and neighboring 250-meter grid cells, consistent with consumption externalities between locations. In contrast, we find no evidence of differences in pre-trends in either group of grid cells beforehand.

We next develop a tractable quantitative model of travel itineraries to rationalize these empirical findings. We consider a city that consists of a set of discrete locations. Agents choose where to live, where to work, and where to consume non-traded services. Utility is derived from consumption of a traded good, non-traded services, and residential floor space. Agents can separate their residence and workplace, to take advantage of differences in wages and amenities, but they face commuting costs that are increasing in travel time. The traded good is costlessly traded and homogenous. The consumption of non-traded services involves costly travel to where they are supplied. Agents have love of variety preferences for non-traded services, which provides the incentive to consume them from multiple locations. But agents face travel costs that are increasing in travel time and incur an intermediate-stay cost for each intermediate location where they consume non-traded services.

Each day, we assume that an agent chooses an ordered sequence of locations to visit, starting and ending at home. On work days, this ordered sequence must include work. On weekends, we allow for any ordered sequence starting and ending at home. Each agent is able to

consume non-traded services from the locations she chooses to visit. Each day, she observes idiosyncratic preference draws for all possible ordered sequences, and chooses the one that offers the highest utility. In making this choice, she trades off her love of variety for non-traded services against the additional travel and intermediate-stay costs from consuming non-traded services from more locations. Travel costs provide the incentive for agents to include multiple intermediate stays within the same travel itinerary (“trip chains”), rather than making separate direct trips between home and each destination. Intermediate-stay costs provide the reason why agents in general choose to visit a strict subset of locations. We show that these travel itineraries give rise to systematic departures from a gravity equation for consumption travel. The probability that an agent travels between an origin and destination depends not only on the travel time between that origin and destination, but also on the travel time of the destination from home and work.

We overcome the high-dimensionality of the choice set implied by travel itineraries using importance sampling. This method uses a Monte Carlo simulation from an auxiliary distribution, and adjusts the sampling rate based on the likelihood ratio between the true distribution and the auxiliary distribution. We choose this auxiliary distribution so that we avoid computing travel probabilities over the set of all possible travel itineraries, thereby avoiding the curse of dimensionality. We use importance sampling to solve the model for a given set of parameter values. We estimate the model’s parameters using a method-of-moments estimation procedure, which minimizes the distance between moments in the simulation and the observed data. We demonstrate that our estimated model provides a good fit to observed patterns of spatial mobility for both targeted and untargeted moments.

We show that our model of travel itineraries is quantitatively successful in rationalizing our quasi-experimental evidence of consumption externalities from the shift to WFH. We interpret the shift to WFH in the model as a change in the commuting technology that allows workers to commute into the office for fewer days each week. In our model of travel itineraries, the resulting reduction in work trips to the central city leads to a collapse in non-work trips there, as workers no longer stop off to consume non-traded services along the way to and from work. In contrast, in a conventional urban model in which all consumption travel occurs directly from home, we find little evidence of a collapse in non-work trips in the central city, because workers consume non-traded services from home regardless of where they work. As a result, our model of travel itineraries matches the observed changes in the relative prices of floor space and the number of non-traded services varieties across locations following the shift to WFH, whereas the conventional urban model does not.

We next show that consumption externalities from travel itineraries generate a quantitatively important force for the agglomeration of economic activity in the central city. We under-

take a counterfactual in which we shut down trip chaining by assuming that intermediate-stay costs become prohibitive. In this counterfactual, we find that employment in *both* traded and non-traded sectors in the most central locations falls by up to one half. The fall in employment in the nontraded sector occurs because this sector relies on demand from in-commuters through their travel itineraries. The fall in employment in the traded sector occurs because the attractiveness of the most central locations as a workplace depends on the availability of nontraded services along commuters’ travel itineraries. These two forces generate a positive feedback loop between the non-traded and traded sectors, thereby providing a force for agglomeration in the central city. Our results indicate that travel itineraries provide a micro-foundation for a flourishing city center with a wide range of non-traded services supported by a dense transport network, as in Tokyo, London, New York City, or Paris.

Finally, we show that travel itineraries provide a stronger rationale for transport infrastructure targeted toward the central city. To illustrate this point, we use our model to evaluate two specific transport improvements, one within the central city connecting subcenters (Yamanote Line), and another connecting the central city to the suburbs (Chuo Line). Abstracting from travel itineraries leads to an undervaluation of the welfare gains from these transport improvements for two main reasons. First, it underestimates travel cost parameters (because of the assumption that all travel occurs from home). Second, it undercounts travel that occurs within trip chains. Together, these two forces lead to a larger undervaluation of the welfare gains from the urban-circular route of the Yamanote Line (61 percent compared to 68 percent), because it is disproportionately used for trip chains. Therefore, taking travel itineraries into account changes the relative welfare gains from alternative transportation infrastructure interventions, in favor of investing in urban commercial centers.

The remainder of the paper is structured as follows. Section 2 discusses the relationship between our research and the existing literature. Section 3 introduces our data. Section 4 presents empirical evidence on properties of travel itineraries that motivate our theoretical model. Section 5 introduces and estimates our travel itinerary model. Section 6 embeds this travel itinerary specification in a general equilibrium quantitative urban model. Section 7 undertakes our counterfactuals. Section 8 concludes.

2 Related Literature

Our work is related to several lines of existing research. We contribute to the theoretical and empirical literature on the internal structure of cities, including monocentric city models (e.g., [Alonso 1964](#); [Mills 1967](#); [Muth 1969](#)), polycentric city models (e.g., [Fujita and Ogawa 1982](#); [Lucas and Rossi-Hansberg 2002](#)), and more recent quantitative urban models (e.g., [Ahlfeldt *et al.*](#)

2015; Allen *et al.* 2017; Monte *et al.* 2018; Dingel and Tintelnot 2023; Tsivanidis 2025). Whereas this literature has traditionally focused on commuting, we develop a tractable model of travel itineraries that can account for the rich patterns of spatial mobility observed in smartphone data, including intermediate stays between home and work.

Other research has argued that endogenous amenities are a force for agglomeration (e.g., Glaeser *et al.* 2001; Florida 2009; Diamond 2016; Gechter and Tsivanidis 2023; Leonardi and Moretti 2023). One strand of this literature studies how endogenous amenities affect residential income segregation, gentrification and demographic sorting (e.g., Allen *et al.* 2017; Balboni *et al.* 2021; Hoelzlein 2023; Couture *et al.* 2024; Almagro and Domínguez-Iino 2025). Another strand analyses the role of spatial and social frictions within cities (e.g., Couture 2016; Davis *et al.* 2019; Su 2022; Hausman *et al.* 2023; Lee and Tan 2024; Vitali 2025). In contrast, we develop a model in which agents consume non-traded services along travel itineraries through geographic space. We show that these travel itineraries provide microeconomic foundations for consumption externalities as a force for agglomeration.

We also contribute to existing research on shopping externalities between stores. Theoretical studies have considered models with stylized geographies, in which shopping externalities provide a reason for stores to colocate, including Eaton and Lipsey (1982); Claycombe (1991); Matsuyama (1995); and Ushchev *et al.* (2015). Empirical research has provided quasi-experimental evidence of consumption externalities from store openings or closings and on-line shopping, including Shoag and Veuger (2018); Benmelech *et al.* (2019); Koster *et al.* (2019); Relihan (2022); and Qian *et al.* (2024). We develop a tractable quantitative model of travel itineraries that microfound these consumption externalities in primitive assumptions about travel behavior and can be used to study their general equilibrium implications. We show that our framework is successful in explaining patterns of spatial mobility in our smartphone data and the collapse in foot traffic in downtown areas following the shift to WFH.

Our paper is also related to empirical research on travel behavior using data from travel surveys and other sources. This research has established a number of stylized facts about travel behavior. First, non-commuting trips are pervasive. For example, in the 2017 National Household Transportation Survey, more than 67 percent of all trips by privately-owned vehicles were for purposes like shopping, errands or recreation (see also CATS 1956, Transportation Research Board 2006 and Agarwal *et al.* 2020).

Second, non-commuting trips are closely related to the availability of nontraded services. Using credit card transactions data for the United States, Einav *et al.* (2021) finds that the extensive margin of the number of customers accounts for around 80 percent of sales variation across establishments, highlighting the importance of local foot traffic. Third, non-commuting trips are concentrated closer to home than commuting trips. Using travel survey and Google

business data for the United States, [Couture \(2016\)](#) finds that the median trip to a restaurant is about 3 miles and lasts 10 minutes, with higher mean values of 6 miles and 14.5 minutes, both of which are shorter than commuting distances and times.

Fourth, non-commuting trips frequently occur as part of a travel itinerary with multiple intermediate stays within a single journey (trip chains), as reviewed in [Thill and Thomas \(1987\)](#). Fifth, the COVID-19 pandemic and the shift to WFH led to substantial changes in spatial mobility. Using U.S. data, [Couture et al. \(2022\)](#) shows that mobility between counties fell substantially during the COVID-19 pandemic. Building on the [Dingel and Neiman \(2020\)](#) measure of the potential for remote working, [Althoff et al. \(2022\)](#), [Delventhal and Parkhomenko \(2022\)](#), [Gokan et al. \(2022\)](#) and [Gupta et al. \(2022b\)](#) show that the centers of large cities were most exposed to the shift to WFH. Comparing small and large cities, [Monte et al. \(2025\)](#) argue that greater WFH in large cities reflects a shift between multiple equilibria. More generally, [Barrero et al. \(2023\)](#) review existing empirical evidence on WFH.

Subsequent to our work, [Oh and Seo \(2023\)](#) develops a model of trip-chaining. Agents face a sequential probabilistic decision problem, such that when choosing whether to visit one location, they do not know whether they will continue to another location or return home, and in which returning home is assumed to be costless. These assumptions imply that the choice probability of traveling to a subsequent destination is the same for all agents and does not depend on their home or work locations. While these assumptions reduce the dimensionality of the state space, they abstract from important real-world features of spatial mobility that are captured by our framework. In particular, the assumption that it is costless to return home is hard to reconcile with evidence that the probability of travelling between an origin and destination is correlated with the travel time of the destination from home and work, even after controlling for the travel time between the origin and destination. In contrast, we allow agents to plan travel itineraries in advance, including the number of destinations and the sequence in which to visit them, anticipating the costs of returning home. Therefore, our model provides a natural explanation for why travel time from home and work matters for travel probabilities, even after controlling for travel times between an origin and destination.

Methodologically, we overcome the high-dimensionality of the choice set of travel itineraries using importance sampling methods following [Kloek and van Dijk \(1978\)](#) and [Ackerberg \(2009\)](#). While we apply our method to travel itinerary choice, it is broadly applicable to discrete choice models with high-dimensional state spaces, such as bundled products and store choices (e.g., [Thomassen et al. 2017](#)), plant location choices (e.g., [Jia 2008](#); [Oberfield et al. 2024](#)), import choice problems (e.g., [Antràs et al. 2014](#)), or route-choice problems for trucking and shipping (e.g., [Brancaccio et al. 2020](#); [Allen et al. 2021a](#); [Allen and Arkolakis 2023](#)).

In each of these settings, existing research either makes additional assumptions, such as

global supermodularity or submodularity (e.g., [Jia 2008](#); [Antràs et al. 2014](#); [Arkolakis et al. 2023](#)); relies on specific functional form assumptions for trade costs (e.g., [Allen and Arkolakis 2023](#)); or limiting parameter values (e.g., [Oberfield et al. 2024](#)). In comparison, our method is typically more computationally intensive, because it involves Monte Carlo simulation.² But an advantage is that we are not required to consider these limiting specifications.

In our framework, locations can be either substitutes or complements for one another in consumption. Therefore, our work relates to existing research in industrial organization that allows choices to be either substitutes or complements, including [Gentzkow \(2007\)](#), [Berry et al. \(2014\)](#), [Thomassen et al. \(2017\)](#), and [Ruiz et al. \(2020\)](#). For example, the opening of a shopping center in one part of town can increase consumption trips to locations that are nearby or along the way (complements in consumption), but reduce consumption trips to an existing shopping center in another part of town (substitutes in consumption).

Finally, our work is related to the growing body of research that uses smartphones or other “big data” to measure spatial mobility. Researchers have used travel surveys (e.g., [Couture 2016](#); [Couture et al. 2018](#), [Zárata 2021](#)); credit card data (e.g., [Agarwal et al. 2020](#); [Allen et al. 2021b](#); [Dolfen et al. 2023](#)); ride-sharing data (e.g., [Gorback 2021](#); [Buchholz et al. 2025](#)); car navigation data (e.g., [Hausman et al. 2023](#)); and cellphone and smartphone data (e.g., [Büchel et al. 2020](#); [Athey et al. 2021](#); [Gupta et al. 2022a](#); [Atkin et al. 2022](#); [Couture et al. 2022](#); [Kreindler and Miyauchi 2023](#)). Our main contribution is to develop a tractable quantitative framework for modelling the rich patterns of spatial mobility observed in these data. We show that travel itineraries give rise to consumption externalities that rationalize our quasi-experimental empirical findings for the shift to WFH, provide a force for agglomeration, and shape the impact of public policy interventions.

3 Data Description

In this section, we introduce our smartphone data and other data sources.

3.1 Smartphone GPS Data

Our main data source is one of the leading smartphone mapping and navigation applications in Japan: *Docomo Chizu NAVI* and *MyDaiz*. Upon installing this application, individuals are asked to give permission to share location information in an anonymized form. Conditional on this permission being given, the application collects the Geographical Positioning System

²While we use Monte Carlo simulation based on importance sampling to overcome the high-dimensionality of travel itinerary decisions, [Dingel and Tintelnot \(2023\)](#) uses Monte Carlo simulation to address small sample variation (granularity) with spatially-disaggregated data. As part of our importance sampling approach, we allow in our Monte Carlo for this small sample variation (granularity).

(GPS) coordinates of each smartphone device up to every 5 minutes (at the highest frequency) whenever the device is turned on (regardless of whether the application is being used). These “big data” provide an immense volume of high-frequency and spatially-disaggregated information on the geographical movements of users throughout each day. The high frequency nature of this data is particularly useful for studying travel itineraries.

The raw unstructured geo-coordinates are pre-processed by the cell phone operator: NTT Docomo Inc. to construct measures of “stays,” which correspond to distinct geographical locations visited by a user during a day. In particular, a stay corresponds to the set of geo-coordinates of a given user that are contiguous in time, whose first and last data points are more than 15 minutes apart, and whose geo-coordinates are all within 100 meters from the centroid of these points.³ The sample size is large; for July 2019 alone, the data include 351 million stays among 3.1 million users (approximately 2.5 percent of the Japanese population). We use data on the sequences of stays of an anonymized random 10% of users with the necessary level of spatial aggregation to deidentify individuals.

This pre-processing also categorizes all stays in each month into three categories of home, work and non-work. “Home” and “work” are defined as the centroids of the first and second most frequent locations of geographically contiguous stays, respectively.⁴ Due to the increased coverage of users over time, we weight each observation by the sampling rate relative to the official residential population for each municipality and year from the population registry of Japan ([Statistics Bureau of Japan 2023](#)). In Subsection 3.2 and Online Appendix A.3.1, we show that our measures of commuting patterns based on our smartphone data closely replicate administrative data on employment by workplace both before and after the COVID-19 pandemic, and bilateral commuting flows. Stays that are neither “home” nor “work” are classified as “non-work” and are interpreted in the model as trips to consume non-traded services.⁵

We focus on the sample of users with home and work locations in the Tokyo Metropolitan Area (which includes the four prefectures of Tokyo, Chiba, Kanagawa and Saitama). Together, these prefectures cover an area of about 13,500 square kilometers and include around 36 million residents. These prefectures are disaggregated into 242 municipalities (excluding islands) and 10,170 Oaza (sub-districts). Using more disaggregated units provides greater spatial reso-

³See Patent Number “JP 2013-89173 A” and “JP 2013-210969 A 2013.10.10” for the detailed proprietary algorithm. This algorithm involves processes to offset the potential noise in measuring GPS coordinates.

⁴See the above patents for the detailed proprietary algorithm. This algorithm assigns work as missing if the user is not active for a sufficient number of days in the data set (among several additional criteria), which applies to about 30 percent of all users. Online Appendix A.8 shows that the probability of work being missing is uncorrelated with the observable characteristics of the municipality of residence.

⁵While some non-work stays could reflect social or business trips, Online Appendices A.3.2, A.5 and A.6.2 provide evidence that shopping trips are the most frequent category of non-work stay and that non-work stays are strongly positively correlated with employment in non-traded service sectors.

lution, but it increases the dimensionality of the state space of travel sequences in the model. We choose a point on this trade-off that achieves reasonable spatial resolution and a manageable state space. Our baseline sample for our travel itineraries model uses Oaza in the 23 wards (municipalities) in Central Tokyo and municipalities in the suburbs, which yields 1,100 spatial units.⁶ In some of our reduced-form specifications, we also report results using Oaza throughout Tokyo or further disaggregating Oaza into 250×250 meter grid cells.

Our baseline estimation sample includes data from the year 2019 before the COVID-19 pandemic. When we examine the shift to WFH in the aftermath of the COVID-19 pandemic, we use data from 2020 to September 2024. When we provide evidence on retail store closures, we also use data from the earlier time period from 2015-9. To abstract from overnight trips, we focus on the sub-sample of user-day observations for which the first and last stay of a travel day is the home location. Travel days are defined to start at 4.00am and end at 3.59am, in line with patterns of night-time home stays in our smartphone data. We refer to sequences of stays that start and end at home on a given travel day as “travel itineraries.” We refer to travel itineraries that include multiple intermediate stays as “trip chains.”

Choosing the time threshold at which to measure a “stay” involves a trade-off. On the one hand, true visits to destinations can be brief (e.g., stopping to pick up a take-away coffee), which implies that a long time threshold could exclude some true stays. On the other hand, travellers can experience delays (e.g., a bus stuck in traffic), which implies that a short time threshold could include some false stays. To the extent that some true visits to locations last less than our threshold of 15 minutes, we could undercount the number of stays, and hence the number of trip chains. Even if such undercounting exists, we show below that trip chains based on our measures of stays are pervasive. Although the pre-processing of our data makes it hard for us to consider thresholds of less than 15 minutes, we report the robustness of our travel itinerary statistics to the use of longer time thresholds of 20 and 30 minutes to measure stays in Online Appendix A.7. We demonstrate a similar pattern of empirical results using these alternative time thresholds, confirming that our empirical findings are not sensitive to the exact time threshold chosen.

We complement our smartphone data with a number of other spatially-disaggregated data sources. We use employment by workplace from the Economic Census (2016, 2021) and commuting flows from the Population Census (2010, 2015) for validating our smartphone data, complementing our descriptive and reduced-form analyses, and calibrating our model. We construct travel times by least-cost path between the centroids of each of our locations us-

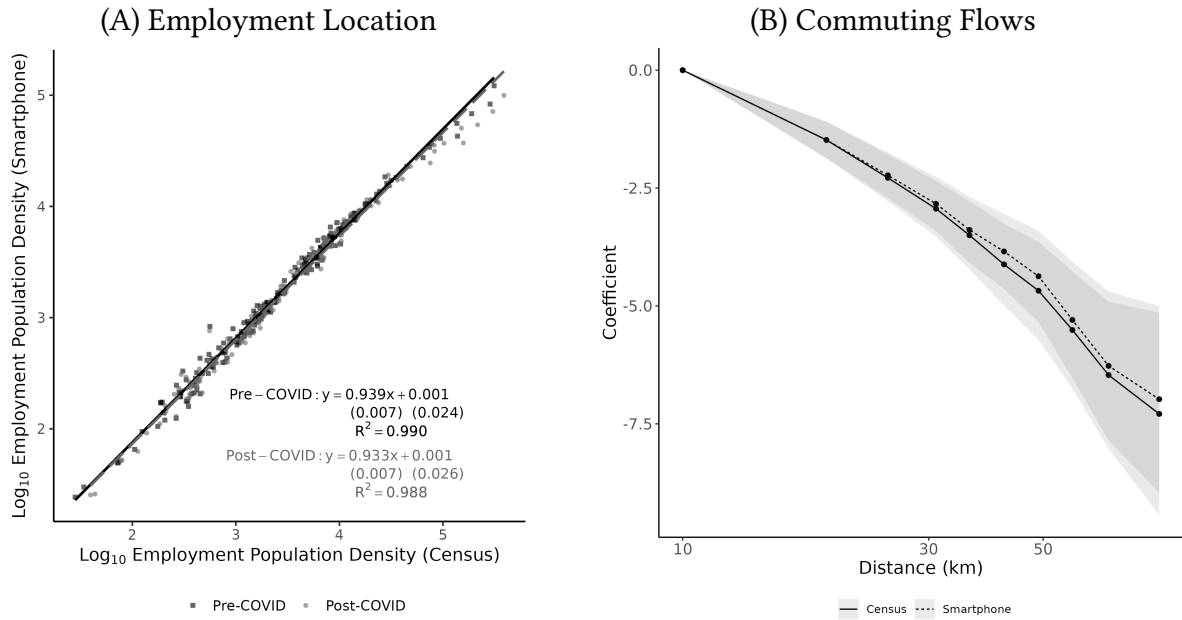
⁶Since the suburbs have lower economic activity than the center, this choice of spatial units helps to ensure more similar levels of economic activity across locations. In Online Appendix A.7, we provide evidence on the robustness of our findings for the properties of travel itineraries across different levels of spatial aggregation.

ing the ArcGIS Network Analyst, shapefiles of railways and bus networks, and assumed travel speeds. We compile a range of data on the shift to WFH in the aftermath of the COVID-19 pandemic, including office rents and the number of establishments obtained from phone directory records. See Appendix A.9 for further details.

3.2 Validation of Smartphone Data Using Census Commuting Data

We now report a validation of our smartphone data and “home” and “work” location classification. In Panel (A) of Figure 1, we show the log density of workers in each Tokyo municipality in our smartphone data against log employment density by workplace in the census data, both for the pre-COVID period (2019) and the post-COVID period (2021). We find a tight and approximately log-linear relationship between them, both in the pre-COVID period (a slope coefficient of 0.939 with an R-squared of 0.990) and the post-COVID period (a slope coefficient of 0.933 with a standard error of 0.007 with an R-squared of 0.988).

Figure 1: Validation of our Smartphone Employment and Commuting Data



Note: In Panel (A), the vertical axis is the log of the number of smartphone users with a work location in the municipality divided by its geographic area (weighted by the sampling rate relative to the official residential population for each municipality and year from the population registry of Japan (Statistics Bureau of Japan 2023)) in 2019 (pre-COVID) and 2021 (post-COVID), and the horizontal axis is the log of employment by workplace in that municipality from the Employment Census in 2016 (pre-COVID) and 2021 (post-COVID) divided by its geographic area. In Panel (B), we report the results of gravity equation estimation including indicator variables for deciles of bilateral distance and residence and workplace fixed effects using the Poisson Pseudo Maximum Likelihood (PPML) estimator; the solid black line and dark gray shading show the estimated coefficients on the decile indicators and 95 percent confidence intervals using the official Population Census data for 2015; the dashed black line and light gray shading show the estimated coefficients on the decile indicators and 95 percent confidence intervals using our smartphone data for 2019. The definitions of home and work in the smartphone data are discussed in the text of Subsection 3.1 above.

In Panel (B) of Figure 1, we show that bilateral commuting flows from our smartphone data have a similar rate of spatial decay with geographic distance as the official census data. We estimate a gravity equation for bilateral commuting flows using both our smartphone commuting data and the census commuting data between pairs of municipalities. We use the Poisson Pseudo Maximum Likelihood (PPML) estimator and include indicators for deciles of bilateral distance and residence and workplace fixed effects. The figure displays the estimated coefficients on the decile indicators (black lines) and the 95 percent confidence intervals (gray shading). The solid black line and dark gray shading show results using the Population Census commuting data for 2015. The dashed black line and light gray shading show results using our smartphone data for our baseline sample period for 2019. We find that these two sets of estimates lie close to one another, particularly for commutes of less than 50 kilometers, which account for the vast majority of observed bilateral commuting flows. In Online Appendix A.3.1, we show that the origin and destination fixed effects and residuals estimated using the census data and smartphone data align closely, providing further evidence that our smartphone data closely replicate patterns of commuting in official census data.

In the Online Appendix, we provide a series of further validation checks on our smartphone data. In Online Appendix A.2, we demonstrate its representativeness by examining its coverage by residence characteristics (income, age and distance to city center) and workplace characteristics (employment by industry and distance to city center). In Online Appendix A.3.2, we show that patterns of spatial mobility in our smartphone data are consistent with those in travel survey data for Tokyo. In Online Appendix A.4, we show that we find an intuitive daily pattern of the timing of home and work stays, with home stays tending to occur during nighttime (outside 6am-9pm), and both work and non-work stays tending to occur during the daytime (from 6am-9pm). In Online Appendix A.6, we show that our smartphone data captures the five stylized facts about travel behavior from existing empirical research that were discussed in the related literature section above.

4 Patterns of Spatial Mobility

We now provide evidence on the empirical properties of travel itineraries that motivate our theoretical modelling, including quasi-experimental evidence of consumption externalities.

4.1 Properties of Travel Itineraries

Travel itineraries with intermediate stays are a pervasive feature of spatial mobility. In Panel A of Table 1, we report the number of work stays and non-work stays (outside home) for all travel itineraries. On weekdays, the average smartphone user makes 1.08 work stays, 1.57

non-work stays, and 1.05 travel itineraries.⁷ On weekends, the corresponding averages are 0.43 work stays, 1.84 non-work stays, and 1.04 travel itineraries. Therefore, on both work and non-work days, most stays occur as part of a single travel itinerary (many of which are multi-stay), instead of making direct trips between home and each destination separately.

Table 1: Patterns of Travel Itineraries

(A) Number of Stays (Outside Home) and Travel Itineraries			
Day Type	Work Stays	Non-work Stays	Travel Itineraries
Weekdays	1.08	1.57	1.05
Weekends	0.43	1.84	1.04

(B) Number of Stays (Outside Home) within a Travel Itinerary			
Number of Stays	All (%)	Weekdays (%)	Weekends (%)
1	42	40	47
2	25	25	25
3	14	14	13
4	7	8	6
5+	12	12	8

(C) Patterns of Travel Itineraries including at Least One Work Stay	
Categories	Share (%)
1-HWH	42
2-HNWH	10
3-HWNH	25
4-HNWNH	8
5-Other	15

Note: Panel (A) shows the number of work stays, non-work stays (outside home), and travel itineraries on weekdays and weekends for all travel itineraries; Panel (B) shows the frequency distribution of the number of stays (including both work and non-work stays outside the home) for all travel itineraries; Panel (C) shows the frequency distribution of sequences of stays for the subset of travel itineraries including at least one work stay: (i) Home-Work-Home (HWH); (ii) Home-Nonwork-Work-Home (HNWH); (iii) Home-Work-Nonwork-Home (HWNH); (iv) Home-Nonwork-Work-Nonwork-Home (HNWNH); (v) Other. The final Other category includes sequences with more than one work stay, such as Home-Work-Nonwork-Work-Home, and those with more than one consecutive non-work stay, such as Home-Nonwork-Nonwork-Work-Home.

Travel itineraries differ substantially in the number of intermediate stops. In Panel B of Table 1, we report the shares of all travel itineraries with each number of stays outside home. On weekdays, 42 percent of travel itineraries contain only one stay outside the home, again confirming that direct trips account for a minority of travel. Around 25 percent of weekday

⁷The average number of work stays can be greater than one on weekdays, because workers can leave their workplace during the day and return there later the same day (e.g., after lunch elsewhere).

travel itineraries include two stays outside the home, with the remaining 33 percent containing three or more stays outside the home. The frequency of observing a travel itinerary is decreasing in the number of stays, except for the final aggregated 5+ category.

Travel itineraries also differ substantially in the sequences of work and non-work stays. In Panel C of Table 1, we consider travel itineraries that include work, and report the shares of these travel itineraries with the following five sequences of stays: (i) Home-Work-Home; (ii) Home-Nonwork-Work-Home; (iii) Home-Work-Nonwork-Home; (iv) Home-Nonwork-Work-Nonwork-Home; and (v) Other, which includes sequences with more than one work stay, such as Home-Work-Nonwork-Work-Home, and those with more than one consecutive non-work stay, such as Home-Nonwork-Nonwork-Work-Home. Category (i) corresponds to a direct trip to work. We find that the majority of these travel itineraries including work also include non-work stays, with the most frequent sequences featuring one non-work stay either on the way to work or back home (category (ii) plus (iii)).

An implication of the inclusion of multiple stays within a single trip is interdependencies in travel behavior, because an agent’s decision to travel between a pair of locations depends on the characteristics of all locations included in the same itinerary.⁸ This interdependence gives rise to consumption externalities between locations. We now use two separate sources of quasi-experimental variation to provide evidence of these consumption externalities.

4.2 Working from Home (WFH)

We begin by using the shift to WFH in the aftermath of the COVID-19 pandemic to provide quasi-experimental evidence of consumption externalities. We use our smartphone data at the Oaza (sub-district) level for the Tokyo Metropolitan Area as a whole.

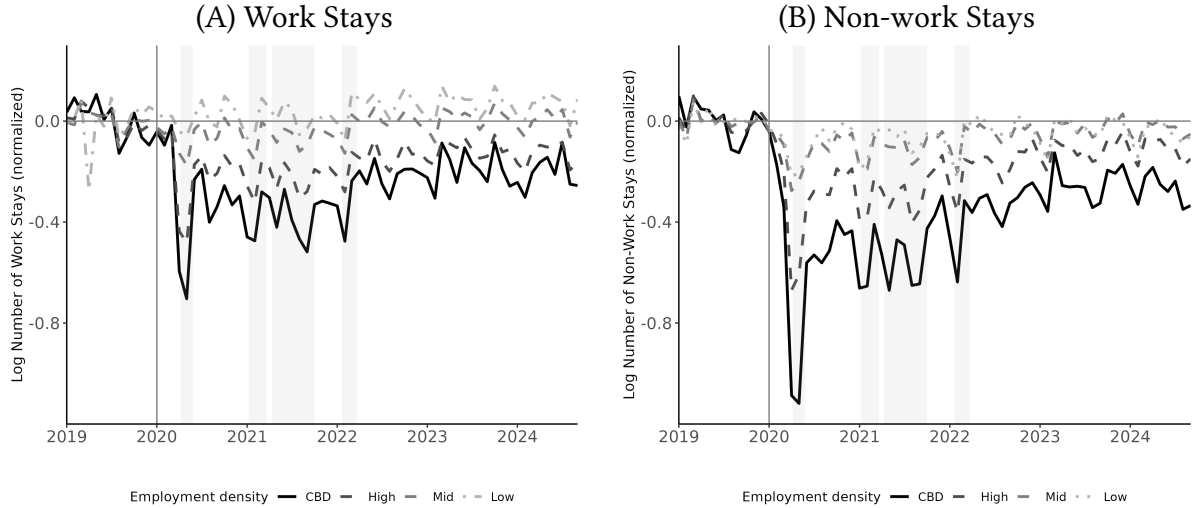
Panel (A) of Figure 2 shows the evolution of mean log work stays from the beginning of 2019 to September 2024. We report these mean levels of foot traffic separately for Oaza with 2km of the central point of Tokyo’s central business district (CBD); for Oaza with high-employment density (top 10 percentiles; excluding those in the CBD), medium-employment density (50-90th percentiles), and low-employment density (0-50th percentiles), where employment density is measured in 2019 before the pandemic. Panel (B) shows the evolution of mean log non-work stays (neither home nor work) over the same time period for the same four groups of locations.

Even after widespread vaccine availability and the end of the COVID-19 pandemic, we find a pronounced reduction in work stays in Tokyo’s CBD and in other high-employment

⁸In Figure A.7.1 of the Online Appendix, we show that non-work stays tend to occur closer to work locations than to home locations during daytime on weekdays, while the opposite is true for nighttime on weekdays and throughout the day on weekends, indicating spatial interdependency of work and non-work stays. We present related evidence of an extended gravity equation and show that our model captures this pattern in Section 5.7.

density locations (Panel A). This pattern of results for work stays is in line with findings for other cities, and is consistent with increased remote and hybrid working, which reduces the number of days on which people commute into work. More notably, we find a similar and quantitatively larger change for non-work stays as for work stays, with the largest reductions observed in Tokyo's CBD and other high-employment density locations (Panel B).

Figure 2: Collapse in Downtown Foot Traffic Following the Shift to Working from Home



Note: Panel (A) shows mean log number of work stays over time; Panel (B) shows mean log number of non-work stays (neither home nor work) over time; gray vertical shading corresponds to periods when the Japanese government issued advisories to stay at home; central business district (CBD) corresponds to within 2km from the centroid of the Oaza of Chiyoda ward; high-density corresponds to Oaza with employment densities in the top 10 percent (excluding those in the CBD); mid-density corresponds to Oaza with employment densities from the 50-90th percentiles (excluding those in the CBD); low-density corresponds to Oaza with employment densities from the 0-50th percentiles (excluding those in the CBD).

The combination of these findings for work and non-work stays reveals the interdependencies in travel behavior and consumption externalities introduced by travel itineraries. As fewer workers commute into work in central locations following the shift to WFH, this reduces the demand for non-traded services (e.g., coffee shops and bars), because fewer people stop off to consume these non-traded services along the way to and from work. The resulting reduction in non-traded employment in central locations further reduces demand for non-traded services, as these displaced workers no longer stop off to consume non-traded services along the way to and from work. Finally, the resulting contraction in available varieties of non-traded services in central locations reduces the attractiveness of commuting into work in these locations. Consistent with these interpretations, we find a drop in office rents and the number of non-traded service establishments in central Tokyo following the shift to WFH (see Online Appendix B including Figures B.4.2 and B.4.3). We show that our estimated travel itinerary model can successfully replicate these patterns in Section 7.1.

Several studies have documented reductions in overall foot traffic in central cities following the shift to WFH (e.g., [Couture et al. 2022](#), [Althoff et al. 2022](#), [Gokan et al. 2022](#), [Delventhal and Parkhomenko 2022](#), [Monte et al. 2025](#); see [Barrero et al. 2023](#) for a survey). Our empirical contribution here is to show that the decline in work trips goes hand-in-hand with a decline in non-work trips, thereby revealing the interdependencies between these two types of travel created by travel itineraries. Our theoretical contribution is to develop a tractable model of travel itineraries that overcomes the resulting high dimensionality of the state space. We embed this tractable model of travel itineraries in a general equilibrium urban model and show that it is quantitatively successful in matching the observed changes in patterns of economic activity following the shift to WFH.

4.3 Closure of Large Retail Stores

We next use the closure of large retail stores to provide a second source of quasi-experimental evidence of consumption externalities. We use data on all large-scale retail stores with a sales floor area of 5,000 square meters or more that were present in the Tokyo metropolitan area from 2015 to 2019. Among these large retail stores, we classify those that had closed by 2019 as the treatment group, and the remaining stores as the control group.

To examine the distance over which consumption externalities extend, we use our smart-phone data disaggregated by 250×250 meter geographical grid cells. We define a grid cell that contains a large retail store as a directly-affected grid cell. We construct a series of concentric distance rings around each directly-affected grid cell to capture spillover effects on neighboring grid cells. Our sample includes the directly-affected and neighboring grid cells for closed large retail stores (treatment group) and surviving large retail stores (control group).

We index large retailers by r and the grid cells containing them by $i(r)$. Observations are grid cells ($i(r)$) and months (t). We denote distance rings by $d \in DR$, where DR is the set of distance rings. We define distance ring 0 as the directly-affected grid cell. The remaining distance rings are of 300 meters width and contain grid cells with centroids up to 2,100 meters from a directly-affected grid cell. For example, the 300-meter distance ring includes grid cells with centroids within 0-300 meters of the centroid of the directly-affected grid cell.⁹

One challenge in estimating the impact of retail store closure is that the decision to close a store is unlikely to be random. Focusing on large stores that are typically part of national chains partially helps to alleviate this concern, because the decision to close a store is often influenced by national considerations, such as the trajectory of sales across the entire na-

⁹To avoid spurious findings of spillovers driven by the directly-affected grid cell of one large retailer lying in the neighboring grid-cell of another large retailer, our baseline specification drops all grid cells in which there is overlap between either the directly-affected and/or neighboring grid cells of multiple large retailers.

tional chain, rather than simply the trajectory of sales for the closed store. We control for time-invariant unobserved heterogeneity between treatment and control grid cells (through the location fixed effect). We also control for differential time trends for directly-affected and neighboring grid cells using distance ring times month fixed effects. Our identifying assumption is parallel trends between the treatment and control grid cells in the absence of the retail store closure. As a check on this identifying assumption, we examine pre-trends in the months leading up to closure. We estimate the following event-study specification:

$$\log Y_{i(r)t} = \sum_{x=-T}^T \sum_{d \in DR} \beta_{dx} (\mathbb{D}_{i(r)}^d \times \mathbb{C}_{i(r)x}) + \eta_{i(r)} + \xi_{dt} + u_{i(r)t} \quad (1)$$

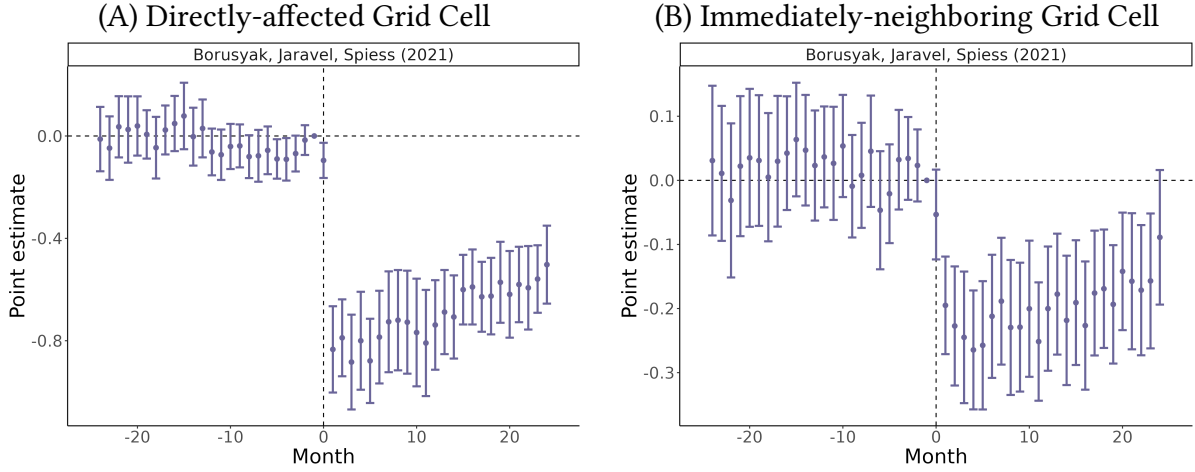
where x denotes months relative to treatment (before or after closure); the excluded category is the month in which closure occurs ($x = 0$); we consider windows of 24 months before and after closure ($T = 24$); β_{dx} are the treatment coefficients of interest on interaction terms between dummy variables for grid cell $i(r)$ lying within distance ring d ($\mathbb{D}_{i(r)}^d$) and dummy variables for grid cell $i(r)$ in period x before or after closure ($\mathbb{C}_{i(r)x}$); $\eta_{i(r)}$ are grid cell fixed effects; ξ_{dt} are distance ring \times month fixed effects; $u_{i(r)t}$ is a stochastic error.

A recent empirical literature has highlighted that the interpretation of the two-way fixed effects estimator in event-study specifications can be problematic in the presence of a variable timing of the treatment (as in our application) and treatment heterogeneity. Therefore, we use the estimator of [Borusyak et al. \(2024\)](#) for our baseline specification, which is robust to this concern. In our empirical application, we find a relatively similar pattern of estimated treatment effects using alternative event-study estimators, as discussed in further detail in Online Appendix C. We weight observations by the number of non-work stays in each grid cell in 2014 before the beginning of our sample period. We cluster the standard errors by grid cell to allow for serial correlation in the error term over time.

Panel (A) of Figure 3 shows the estimated treatment effects for the directly-affected grid cell (β_τ) from our event-study specification using the [Borusyak et al. \(2024\)](#) estimator. We find little evidence of differences in pre-trends, with estimated coefficients for negative treatment months that are close to zero and statistically insignificant. Immediately after the large retail store closure, we find a substantial and statistically significant decline in the number of non-work stays in the treated location by approximately 0.80 log points. This decline persists for at least 24 months, with attenuation over time, consistent with new uses being found for the vacated floor space.

Panel (B) of Figure 3 shows the estimated spillover effects for the immediately-neighboring 300-meter distance ring. As for the directly-affected location, we find little evidence of differences in pre-trends, with estimated coefficients for negative treatment months that are close

Figure 3: Direct and Spillover Effects of the Closure of Large Retail Stores on Non-work Stays



Note: Estimated treatment-month interactions (β_{dx}) for the impact of large retail store closures (floor space $> 5,000$ meters squared) on non-work stays from equation (1); Panel (A) shows estimates for the directly-affected grid cell; Panel (B) shows estimates for the immediately-neighboring 300-meter distance ring; observations are 250×250 meter grid cells and months in Tokyo from January 2015 to December 2019; dependent variable is the log number of non-work stays (neither home nor work); event-study specification estimated using the [Borusyak et al. \(2024\)](#) estimator; treatment months before -24 and after +24 are binned into the first and last categories; circles denote the point estimates; vertical bars indicate the 95 percent confidence intervals based on standard errors clustered by 250×250 meter grid cell; estimation with the full sample is infeasible because of the large control group and hence we report results using a 30 percent random sample of the control group.

to zero and statistically insignificant. Immediately after the large retail store closure, we find a substantial decline in non-work stays in neighboring locations of around 0.20 log points, which is statistically significant at conventional critical values. Again this decline persists for around 24 months, with attenuation over time, consistent with new uses being found for the vacated floor space.¹⁰

Therefore, consumption externalities are not only a feature of the shift to WFH, but are also observed more broadly for the closure of large retail stores. Travel itineraries provide a natural explanation for these empirical findings. As the retail store closure reduces foot traffic in the location containing the closed store, this in turn depresses foot traffic in neighboring locations, because fewer agents now stop off in those neighboring locations as convenient intermediate stops along travel itineraries to the closed store.

¹⁰These findings of consumption externalities using our Japanese smartphone data are consistent with a range of existing evidence for retail store openings and closings, including [Shoag and Veuger \(2018\)](#), [Benmelech et al. \(2019\)](#), [Koster et al. \(2019\)](#), [Relihan \(2022\)](#) and [Qian et al. \(2024\)](#). In Online Appendix C, we provide further evidence on the spatial decay of these consumption externalities.

5 Travel Itinerary Model

In this section, we first develop our model of travel itineraries. We next estimate this model using our smartphone data and show that it provides a good approximation to observed patterns of spatial mobility in our smartphone data.

5.1 Theoretical Framework

We assume that preferences are separable in the consumption of traded goods and non-traded services, which enables us to focus in this section on non-traded services consumption. We assume that the sub-utility function for non-traded services is homothetic, which allows us to characterize consumption choices using the unit expenditure function or price index.

We consider a city that consists of a set of locations: $\mathbb{N} \equiv \{1, \dots, \mathcal{N}\}$. We model the choice of travel itinerary for an agent with home $h \in \mathbb{N}$ and workplace $j \in \{\mathbb{N}, \emptyset\}$ on a given day. The inclusion of $j = \emptyset$ in the set of workplaces allows for non-workdays (weekends) when an agent does not travel to work. We index travel itineraries by $I \in \mathcal{I}_{hj}$, where \mathcal{I}_{hj} is the set of possible travel itineraries for an agent with home h and workplace j . We denote $C(I) \subseteq \mathbb{N}$ as the subset of locations at which non-traded services are consumed in travel itinerary I . Travel itinerary I corresponds to an ordered sequence of $C(I)$ that starts and ends at h and includes j (allowing $j = \emptyset$ on non-workdays).¹¹ Therefore, the set of possible travel itineraries for each residence h and workplace j (\mathcal{I}_{hj}) includes all ordered subsets of locations that satisfy these conditions (e.g., $\{h, 1, j, 14, h\}$, $\{h, j, 25, h\}$, where 1, 14 and 25 are specific locations). We assume that agents make only one travel itinerary each day, which approximates the patterns that we observe in our smartphone data in Table 1.

We consider the subutility from consuming non-traded services for agent ω with home h and workplace j and travel itinerary I ($U_{\omega I|hj}^S$). This subutility depends on the cost of the non-traded services consumed at each location $i \in C(I)$ in that itinerary inclusive of travel costs ($\mathbb{P}_{I|hj}$) and an idiosyncratic preference shock ($\epsilon_{\omega I}$):

$$U_{\omega I|hj}^S = \frac{\epsilon_{\omega I}}{\mathbb{P}_{I|hj}}. \quad (2)$$

The cost of non-traded services inclusive of travel costs ($\mathbb{P}_{I|hj}$) depends on the non-traded price indexes $\{P_{nS}\}$ at each location $n \in C(I)$ and travel costs for the itinerary ($\tau_{I|hj}$):

$$\mathbb{P}_{I|hj} = g_{I|hj}(\{P_{nS}\}_{n \in C(I)}, \tau_{I|hj}). \quad (3)$$

The subscript ($I|hj$) on the cost function ($g_{I|hj}(\cdot, \cdot)$) indicates that the cost of consuming non-traded services depends on home h , workplace j , and the itinerary I (through travel costs).

¹¹We allow $C(I)$ to include the home and/or work locations, in which case the agent chooses to consume nontraded services in the home and/or work locations (such that $C(I)$ includes h and/or j).

We do not impose a particular parametric assumption on this cost function until we discuss our quantitative implementation in Section 5.3.

The idiosyncratic preference shock ($\epsilon_{\omega I}$) captures all the idiosyncratic factors that can lead an agent to choose to consume non-traded services from a particular ordered sequence of locations, such as preferences for specific routes. We assume for simplicity that these idiosyncratic preferences are drawn independently from a Fréchet distribution:

$$F(\epsilon) = e^{-\epsilon^{-\theta}}, \quad \theta > 1, \quad (4)$$

where the shape parameter θ controls the dispersion of idiosyncratic preferences, and hence the responsiveness of travel itinerary choices to costs relative to idiosyncratic factors. We set the scale parameter to one, but allow locations to differ in their productivity in supplying non-traded services and bilateral travel costs.

Although the independence assumption is strong, it ensures that the itinerary choice problem remains tractable, despite our high-dimensional choice space. This assumption is a common benchmark in the literature on transport network routing (Allen and Arkolakis 2023) and truck routing (Allen *et al.* 2021a). Whereas the existing transport literature typically considers one-way shipments among a fixed set of locations, two important features of personal travel itineraries are that agents choose to visit an endogenous set of locations (as well as an endogenous path between them) and return home at the end of the day. As a result, travel itineraries introduce complex interdependencies into agents' travel decisions. The decision to travel between any pair of locations depends on the characteristics of other locations included in the same travel itinerary. Some pairs of locations can be complements in agents' travel decisions, while other pairs of locations can be substitutes.¹²

Each agent ω with home h and workplace j chooses her travel itinerary I_ω on a given day to maximize her subutility from consuming non-traded services:

$$I_\omega = \max_{\ell \in \mathcal{I}_{hj}} \left\{ \frac{\epsilon_{\omega \ell}}{\mathbb{P}_{\ell|hj}} \right\}. \quad (5)$$

Using the properties of the Fréchet distribution, the probability that the agent chooses travel itinerary I depends on the relative costs ($\mathbb{P}_{\ell|hj}$) of all possible travel itineraries $\ell \in \mathcal{I}_{hj}$:

$$\Lambda_{I|hj} = \frac{\mathbb{P}_{I|hj}^{-\theta}}{\sum_{\ell \in \mathcal{I}_{hj}} \mathbb{P}_{\ell|hj}^{-\theta}}. \quad (6)$$

¹²The function $g_{I|hj}(\cdot, \cdot)$ in equation (3) allows for arbitrary aggregation of location price indices, which captures complementarity and substitution between locations. In Online Appendix D, we provide first-order comparative statics for the effect of changes in the non-traded services price index in a given location on non-work stays in surrounding locations, and show that our model captures complementarity and substitution.

Note that the set of possible travel itineraries given home h and workplace j (\mathcal{I}_{hj}) is potentially extremely high-dimensional. With \mathcal{N} locations, there are $2^{\mathcal{N}}$ combinations of locations that an agent can visit, and many more sequences in which to visit them. Therefore, in empirical settings using spatially-disaggregated data (e.g., $\mathcal{N} = 1, 100$ in our application), estimation using conventional methods quickly becomes infeasible. We show below how we overcome this curse of dimensionality using importance sampling.

The expected utility from consuming non-traded services per unit of expenditure across the set of possible travel itineraries (\mathcal{I}_{hj}) for a given home h and workplace j involves taking the expectation across the distribution for idiosyncratic preferences:

$$\mathbb{E}_{\epsilon} \left[\frac{\epsilon_{\omega I}}{\mathbb{P}_{I|hj}} \right] = \mathbb{A}_{hj}, \quad \mathbb{A}_{hj} \equiv \varrho \left[\sum_{\ell \in \mathcal{I}_{hj}} \mathbb{P}_{\ell|hj}^{-\theta} \right]^{\frac{1}{\theta}}, \quad (7)$$

where $\mathbb{E}_{\epsilon} [\cdot]$ denotes the expectation with respect to the distribution of ϵ ; and $\varrho \equiv \Gamma \left(\frac{\theta-1}{\theta} \right)$, where $\Gamma(\cdot)$ is the Gamma function.

We refer to \mathbb{A}_{hj} as consumption access for non-traded services, because it summarizes the expected utility from travel itineraries to consume non-traded services net of travel costs, and corresponds to the inverse of the expected cost of these travel itineraries to consume non-traded services. An increase in the number of locations (an increase in the cardinality of the set \mathcal{I}_{hj}) raises consumption access, because of horizontal differentiation across locations. An increase in the cost of a given itinerary $\ell \in \mathcal{I}_{hj}$ ($\mathbb{P}_{\ell|hj}$) reduces consumption access, because it increases the expected expenditure required to obtain a given level of expected utility from travel itineraries to consume non-traded services.

Our measure of consumption access for non-traded services (\mathbb{A}_{hj}) differs in important ways from market access measures in conventional urban models. In [Ahlfeldt *et al.* \(2015\)](#) and [Tsvanidis \(2025\)](#), there are no costs to trading goods, which implies that commuter market access is a sufficient statistic for location decisions. In [Monte *et al.* \(2018\)](#) and [Allen *et al.* \(2017\)](#) there are costs to trading goods, which implies that both goods and commuter market access matter for location decisions. However, all trade costs are incurred bilaterally between home and the place of production. In contrast, in our framework, agents themselves travel through space to consume non-traded services, which creates the incentive to chain destinations along an itinerary to economize on total travel costs. In the partial equilibrium model of restaurant choice of [Davis *et al.* \(2019\)](#), agents choose whether to visit a restaurant from home, work or the path between home and work. But they are not allowed to chain sequences of destinations within a single travel itinerary. In contrast, we develop a tractable model of travel itineraries, embed it in a general equilibrium urban model, and evaluate the quantitative relevance of travel itineraries for the agglomeration of economic activity in urban areas.

5.2 Importance Sampling Method

The main challenge for estimating our model of travel itineraries and undertaking counterfactuals is the high dimensionality of the set of possible travel itineraries (\mathcal{I}_{hj}). Evaluating the probability of observing a travel itinerary ($\Lambda_{I|hj}$) in equation (5) involves computing the sum over the set of possible travel itineraries (\mathcal{I}_{hj}) in the denominator, which quickly becomes impractical in empirical settings with spatially-disaggregated data.

To overcome this challenge, we develop a method to simulate the probability of choosing a travel itinerary ($\Lambda_{I|hj}$) based on the method of importance sampling (Kloek and van Dijk 1978). We use importance sampling in an inner loop to solve the model for a given set of parameter values. We estimate the model's parameters in an outer loop using the simulated method of moments (SMM), as discussed in Section 5.4 below.

The basic idea for our importance sampling approach is as follows. First, we obtain a Monte-Carlo sample of travel itineraries from an auxiliary distribution that can be computed exactly. Second, we adjust the sampling rate from this auxiliary distribution based on the likelihood ratio between the true distribution and the auxiliary distribution. We thereby obtain simulated travel itinerary probabilities ($\Lambda_{I|hj}^*$) that are consistent estimates of the true travel itinerary probabilities ($\Lambda_{I|hj}$). Importantly, we show that these simulated travel itinerary probabilities ($\Lambda_{I|hj}^*$) can be constructed without having to compute the summation over the set of possible travel itineraries in the denominator of equation (5), because this summation cancels from the likelihood ratio weights. Formally, the procedure is defined in Algorithm 1.

Algorithm 1 (Importance Sampling) Denote the auxiliary probability distribution of travel itineraries by agents with home h and workplace j by $F_{hj}(I)$, defined over the set of possible travel itineraries \mathcal{I}_{hj} . The simulated probability $\Lambda_{I|hj}^*$ is defined as follows:

1. Draw R itineraries $\{I_r\}$ from the auxiliary distribution $F_{hj}(\cdot)$. Denote the empirical distribution of the simulated draws by $\mathcal{E}_{I|hj} = \frac{1}{R} \sum_{1, \dots, R} 1[I = I_r]$.
2. Weight each draw by the likelihood ratio $\Lambda_{I|hj}/F_{hj}(I) \propto \mathbb{P}_{I|hj}^{-\theta}/F_{hj}(I)$ to obtain the simulated probability distribution $\Lambda_{I|hj}^*$:

$$\Lambda_{I|hj}^* = \frac{\mathcal{E}_{I|hj} \Lambda_{I|hj}/F_{hj}(I)}{\sum_{\ell \in \mathcal{I}_{hj}^R} \mathcal{E}_{\ell|hj} \Lambda_{\ell|hj}/F_{hj}(\ell)} = \frac{\mathcal{E}_{I|hj} \mathbb{P}_{I|hj}^{-\theta}/F_{hj}(I)}{\sum_{\ell \in \mathcal{I}_{hj}^R} \mathcal{E}_{\ell|hj} \mathbb{P}_{\ell|hj}^{-\theta}/F_{hj}(\ell)}, \quad (8)$$

where \mathcal{I}_{hj}^R is the subset of \mathcal{I}_{hj} that are sampled in Step 1, i.e., $\mathcal{E}_{I|hj} > 0$.

The simulated probability distribution in equation (8) has an intuitive interpretation. Compared to the actual travel itinerary choice probability $\Lambda_{I|hj}$, the Monte-Carlo draws from $F_{hj}(\cdot)$

under-sample itineraries with higher likelihood ratios $\Lambda_{I|h_j}/F_{h_j}(I)$. Therefore, re-weighting each draw by this likelihood ratio yields a consistent estimator of $\Lambda_{I|h_j}$.

This likelihood ratio $\Lambda_{I|h_j}/F_{h_j}(I)$ is proportional to $\mathbb{P}_{I|h_j}^{-\theta}/F_{h_j}(I)$, which omits the summation over the set of possible itineraries from the denominator of the travel itinerary probabilities ($\Lambda_{I|h_j}$) in equation (5). This summation is a normalizing constant that cancels from the numerator and denominator in equation (8). By abstracting from this normalizing constant, we avoid having to directly compute the denominator of $\Lambda_{I|h_j}$ that is subject to the curse of dimensionality.¹³ We still need to compute the cost of consuming non-traded services for a given itinerary ($\mathbb{P}_{I|h_j}$), but this object is of much smaller dimension, involving a summation over the destinations included in that travel itinerary.

An important advantage of this algorithm is that the choice of the auxiliary distribution $F_{h_j}(\cdot)$ does not affect the results asymptotically as $R \rightarrow \infty$. This property can be seen from equation (8), in which $\Lambda_{I|h_j}^* \rightarrow \Lambda_{I|h_j}$ as $R \rightarrow \infty$, as long as the support of $F_{h_j}(\cdot)$ has common support with $\Lambda_{I|h_j}$.¹⁴ Away from this asymptotic limit, for finite values of R , the precision of the approximation depends on how close $F_{h_j}(\cdot)$ is to the original distribution $\Lambda_{I|h_j}$, as discussed in Kloek and van Dijk (1978) and Akerberg (2009).¹⁵

In our empirical application, we find that the following choice of auxiliary distribution $F_{h_j}(\cdot)$ performs well in practice. First, we assume that each agent randomly draws a number of non-work stays at which to consume non-traded services along a travel itinerary that day. Second, we assume that each agent sequentially picks destinations for each non-work stay myopically, without taking into account the continuation value of visiting a destination. Appendix E.1 formally describes this choice of auxiliary distribution in further detail. Appendix E.2 shows that in practice, the sampling errors are negligible for our quantitative results given our choice of the number of importance samples $R = 200$ per home and work.

We also use importance sampling to overcome the challenge of measuring consumption access (\mathbb{A}_{h_j}). For the same reason that computing the travel itinerary probabilities ($\Lambda_{I|h_j}$) in equation (5) is difficult, computing consumption access (\mathbb{A}_{h_j}) in equation (7) is also problematic, because it involves a summation over the set of all possible travel itineraries (\mathcal{I}_{h_j}). Instead of directly computing this summation, we overcome this challenge by using equations (5) and

¹³Even with a single non-work stay (on either the way to or from work), evaluating this summation in the denominator is computationally costly, because it needs to be evaluated for each pair of home and work.

¹⁴Note that $\mathcal{E}_{I|h_j} \rightarrow F_{h_j}(I)$ as $R \rightarrow \infty$. Therefore, as long as $F_{h_j}(I) > 0$ whenever $\Lambda_{I|h_j} > 0$, we have $\Lambda_{I|h_j}^* \rightarrow \Lambda_{I|h_j}$ in equation (8).

¹⁵An additional computational advantage of importance sampling is that researchers can use the same draws from the auxiliary distribution for different parameter values, and simply update the likelihood ratio for these different parameter values during estimation and counterfactual analysis, as discussed in Akerberg (2009).

(7) to re-write consumption access (\mathbb{A}_{hj}) in terms of the travel itinerary probabilities:

$$\mathbb{A}_{hj} = \varrho \left[\frac{\mathbb{P}_{I|hj}^{-\theta}}{\Lambda_{I|hj}} \right]^{\frac{1}{\theta}}, \quad (9)$$

for any given travel itinerary I . We replace the true travel itinerary probabilities ($\Lambda_{I|hj}$) with the simulated probabilities from our importance sampling ($\Lambda_{I|hj}^*$), using the travel itinerary I that is most frequently drawn from our auxiliary distribution for home h and work j .

5.3 Model Parametrization

We now parameterize the non-traded services cost function ($g_{I|hj}(\cdot, \cdot)$) in equation (3) by assuming that agents have constant elasticity of substitution (CES) preferences over the horizontally-differentiated varieties supplied by each location:

$$\mathbb{P}_{I|hj} = \left(\sum_{n \in C(I)} P_{nS}^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tau_{I|hj}^{-1}, \quad (10)$$

where recall that P_{nS} is the price index for non-traded services (subscript S stands for services) in a given location n ; $\sigma > 1$ is the elasticity of substitution across non-traded services varieties; the first term is a CES price index across prices of non-traded services in each location; the second term adjusts this CES price index for the total travel costs ($\tau_{I|hj}$) of the itinerary.

We assume that the total travel cost for itinerary I ($\tau_{I|hj}$) depends multiplicatively on the travel costs for each leg $m \in \{1, \dots, M(I)\}$ included in that itinerary ($\tilde{\tau}_m \geq 1$) as follows:

$$\tau_{I|hj} = \varsigma_{hj} \left(\prod_{m=1}^{M(I)} \tilde{\tau}_m \right), \quad (11)$$

where ς_{hj} is a normalizing parameter specific to each pair of home h and workplace j that is discussed further below.

This specification implies that total travel costs increase at a constant proportional rate with the number of legs $M(I)$ in itinerary I (holding constant travel time per leg) and in travel costs per leg $\tilde{\tau}_m$ (holding constant the number of legs). While, in principle, one could allow a more flexible relationship between total travel costs and travel costs per leg, we find that this multiplicative specification provides a good fit to the data.

We assume that travel costs for each leg m ($\tilde{\tau}_m$) are a constant elasticity function of the travel time between the origin $i(m)$ and destination $k(m)$ for that leg ($T_{i(m),k(m)}$):

$$\tilde{\tau}_m = \begin{cases} \eta_W (T_{i(m),k(m)})^{\rho_W} & \text{if weekday } (j \neq \emptyset) \\ \eta_N (T_{i(m),k(m)})^{\rho_N} & \text{if non-workday } (j = \emptyset) \end{cases} \quad (12)$$

where $T_{i(m),k(m)} > 0$ is the travel time in minutes.

The parameters ρ_W and ρ_N control the elasticities of the travel cost for each leg to the travel time.¹⁶ They capture the idea that travel itineraries involving longer total travel time are more costly. The parameters η_W and η_N regulate the impact on travel costs of adding an additional intermediate stay to consume non-traded services to a travel itinerary. They capture the cost of breaking a journey (e.g., finding parking, entering and exiting a train station, or the psychological cost of interrupting a journey) relative to the benefit of visiting an additional location. We allow both sets of parameters to differ between workdays (superscript W) and non-workdays (superscript N), because the costs of travel time and breaking a journey could be greater during work days.

We set the normalizing parameter specific to each pair of home h and workplace j as equal to $\varsigma_{hj} \equiv (\tilde{\tau}_{hj}\tilde{\tau}_{jh})^{-1}$, such that the travel costs for each itinerary are measured relative to those for an itinerary with no intermediate stays and only direct travel between home and work. This normalizing parameter is a constant for a given home h and workplace j . Therefore, it has no effect on the probability of choosing alternative travel itineraries for a given home and workplace, where home and workplace are determined before daily travel itinerary decisions. We introduce this normalizing parameter here to be consistent with the specification of commuting choices when we embed our travel itinerary model in general equilibrium in Section 6 below. This specification includes a separate term for the costs of travelling between home and work, in order to be able to nest the conventional urban model as the special case of our model with no consumption of non-traded services.

In taking the model to the data, we take into account that our smartphone data are aggregated into spatial units of different sizes. To accommodate this feature, we assume that the idiosyncratic preferences in equation (4) are drawn at the level of disaggregated grid cells within each spatial unit, following the strategy in Kreindler and Miyauchi (2023). Under this assumption, the mean idiosyncratic preference draw for each travel itinerary depends on the total area of the geographical units included in the spatial units in that travel itinerary. We thus obtain the following generalization of the itinerary choice probability (6):

$$\Lambda_{I|hj} = \frac{Z_{I|hj} \mathbb{P}_{I|hj}^{-\theta}}{\sum_{\ell \in \mathcal{I}_{hj}} Z_{\ell|hj} \mathbb{P}_{\ell|hj}^{-\theta}}, \quad (13)$$

where $Z_{I|hj} = \prod_{i \in C(I)} Z_i$ and Z_i is the land area of each spatial unit, as discussed further in Online Appendix F.

From equation (10), the non-traded services supplied by different locations are substitutes for a given set of locations included in each travel itinerary ($C(I)$) with an elasticity of sub-

¹⁶We choose an iso-elastic rather than an iso-semi-elastic travel cost specification, because it better approximates the fat-tailed spatial mobility patterns in our smartphone data (e.g., Figure A.6.3 of the Online Appendix).

stitution $\sigma > 1$. However, this set of locations ($C(I)$) is itself endogenous and responds to changes in the price indexes for non-traded services in each location. Once this endogenous extensive margin is taken into account, the non-traded services supplied by different locations can be either substitutes or complements, as shown formally in Online Appendix D.¹⁷

Under our assumption of constant elasticity of substitution preferences for non-traded services, the price index for non-traded services in location n (P_{nS}) is:

$$P_{nS} = M_n^{-\frac{1}{\sigma-1}} p_n$$

where M_n is the mass of varieties per unit land area in location n ; and p_n is the price of the representative variety supplied by that location. In our general equilibrium model below, we determine these endogenous variables using profit maximization, monopolistic competition, and free entry.

Our baseline model of travel itineraries in this section makes a number of simplifying assumptions, including aggregating all types of non-traded services together, CES preferences, and travel costs for each leg of an itinerary that are a constant elasticity function of travel time. We make these simplifying assumptions in order to demonstrate as clearly as possible the role of travel itineraries in generating interdependence in travel decisions and consumption externalities between locations. Despite these simplifying assumptions, we show that our estimated model provides a good fit to observed patterns of spatial mobility. In contrast, the special case of our model in which all consumption trips occur directly from home is unsuccessful in capturing these observed patterns of spatial mobility.

5.4 Parameter Estimation

We next discuss the estimation of the parameters of our travel itinerary model. We first determine the elasticity of substitution for non-traded services (σ) using data on the Japanese retail sector. Under our assumption of CES preferences and monopolistic competition (consistent with our general equilibrium model in Section 6), the equilibrium ratio of variable profits to revenue for non-traded services in the model is $1/\sigma$. In the Japanese Economic Census, the ratio of profits to revenue in the retail sector is approximately 22 percent, which implies $\sigma = 4.6$. This parameter value is marginally smaller than those used in related settings for the U.S.: [Couture et al. \(2024\)](#) assumes an elasticity of substitution of 6.8 across the non-traded services supplied by neighborhoods; [Hottman et al. \(2016\)](#) estimates a median elasticity of substitution across barcoded goods in the retail sector of 6.9.

¹⁷For example, a fall in the price of non-traded services in one location increases trips to consume non-traded services in that location, which can increase trips to consume non-traded services in neighboring locations that are convenient intermediate stops along the way (complements), and decrease consumption trips to other locations that are not convenient intermediate stops (substitutes).

We estimate the travel parameters of the model ($\Theta \equiv \{\theta, \rho_W, \rho_N, \eta_W, \eta_N\}$) using a simulated method of moments (SMM) procedure and moments of travel itineraries from our smartphone data. We start with an initial guess for the travel parameters Θ , before updating this initial guess to minimize the sum of squared deviations between the moments in the simulated data from the model and the moments in our smartphone data.

Non-traded Price Indexes Given an initial parameter guess (Θ), we invert the model and solve for the non-traded price indexes in each location $\{P_{nS}\}$ that rationalize the observed share of travel destinations for a subset of travel itineraries. Denoting this subset of travel itineraries by $\mathcal{I}_{hj}^{\text{invert}} \subset \mathcal{I}_{hj}$, we solve for the unique set of non-traded price indexes $\{P_{nS}(\Theta)\}_n$ for which the probability of visiting a destination to consume non-traded services in the simulated model (right-hand side) is equal to the observed probability in the data (left-hand side):

$$\sum_{h,j} \sum_{I \in \mathcal{I}_{hj}^{\text{invert}}} S_{I|h,j} \Omega_{hj} = \sum_{h,j} \sum_{I \in \mathcal{I}_{hj}^{\text{invert}}} \Lambda_{I|h,j}(\Theta, \{P_{nS}\}_n) \Omega_{hj}, \quad (14)$$

where Ω_{hj} is the fraction of population with home h and work j averaged across workdays and non-workdays; $S_{I|h,j}$ is the observed choice probability of itinerary I for home h and work j in our smartphone data; and recall that $\Lambda_{I|h,j}(\Theta, \{P_{nS}\}_n)$ is the probability of choosing itinerary I for home h and work j in the model.

We choose as the subset $\mathcal{I}_{hj}^{\text{invert}}$ those travel itineraries with only one non-work stay.¹⁸ This choice ensures a unique solution for the set of non-traded price indexes $\{P_{nS}(\Theta)\}_n$, because the property of connected substitutes is satisfied for these choice probabilities. More generally, for travel itineraries with more than one non-work stay, this price index inversion need not be unique. The reason is that connected substitutes need no longer hold, because of the complementarities between the decisions to consume non-traded services from multiple destinations that are part of the same travel itinerary.

We implement our model inversion for the subset of travel itineraries with only one non-work stay using a standard algorithm (e.g., [Alvarez and Lucas 2007](#)). We guess an initial vector for non-traded price indexes and compute the vector excess demands in equation (14). If all elements of this vector are zero, we have found the unique solution for non-traded price indexes. If not, we update our initial guess by increasing price indexes in locations with positive excess demands, reducing price indexes in locations with negative excess demands, and iterating until convergence. Since this excess demand system satisfies connected substitutes for the subset of travel itineraries with only one non-work stay, this algorithm necessarily converges to the unique solution.

¹⁸This subset of travel itineraries includes (i) Home-Nonwork-Home, (ii) Home-Nonwork-Work-Home, and (iii) Home-Work-Nonwork-Home.

Using our unique solutions for $\{P_{nS}(\Theta)\}_n$ from itineraries with only one non-work stay, we show below that our model provides a good fit to the observed frequencies of travel itineraries with more than one non-work stay. Therefore, although the travel itineraries with only one non-work stay used to solve for non-traded price indexes are a particular type of trip, we find that our model provides enough structure for us to successfully extrapolate to travel itineraries with more than one non-work stay.

Travel Itinerary Moments Given our initial guess for the parameters Θ and solution for the non-traded price indexes $\{P_{nS}(\Theta)\}_n$, we compute the probabilities of travel itineraries in the model $(\Lambda_{I|h_j}(\Theta, \{P_{nS}(\Theta)\}_n))$. Using these simulated travel itinerary probabilities, we construct moments $(m_n(\Theta))$ for each location as the difference between the predictions of our model and the empirical values in the data. We use the following moment conditions:

1. We compute the fraction of travel itineraries involving two non-work stays and those involving three or more non-work stays during workdays and non-workdays. These four moments are informative about the intermediate-stay costs (η_W and η_N).
2. We compute the 10th, 50th, and 90th percentiles of the distribution of travel times between two consecutive stays on workdays and non-workdays. These six moments are informative about the travel time parameters ρ_W and ρ_N .
3. We compute log deviations between the model's predictions for non-traded price indices ($P_{nS}(\Theta)$) and empirical measures of these price indexes for each location n . We compute these empirical price indexes using our assumed CES functional form ($P_{nS} = M_n^{-1/(\sigma-1)} p_n$), observed data on the density of retail establishments per unit area (M_n) and the average retail price (p_n) in each location from the Retail Price Survey ([Statistics Bureau of Japan 2019](#)), and the elasticity of substitution (σ) as determined above. This moment is informative about the dispersion in idiosyncratic preferences for non-traded services (θ).

SMM Estimator We update our initial guess for the travel parameters Θ to minimize the sum of squared deviations between our simulated model's predictions and the empirical values of the moments in the data:

$$\hat{\Theta} = \min_{\Theta} \left\{ \left(\frac{1}{N} \sum_{n=1}^N m_n(\Theta) \right)' \mathbf{W} \left(\frac{1}{N} \sum_{n=1}^N m_n(\Theta) \right) \right\}, \quad (15)$$

where \mathbf{W} is the weighting matrix.

We use the efficient weighting matrix (\mathbf{W}) from a two-step SMM estimation procedure. We compute the standard errors of the estimated parameters using the asymptotic variance-covariance matrix adjusted by the number of simulation draws (McFadden 1989).

Model With Only Direct Consumption Trips In addition to our baseline model, we report results for a special case that corresponds to a version of the conventional urban model, in which all consumption travel occurs through direct trips from home. We assume that travel costs take the same form as in equations (11) and (12), except that the single travel itinerary each day to consume non-traded services consists of a direct trip from home to a consumption location and back again. In this special case, the multiplicative parameters (η_W, η_N) for the cost of each non-work stay cancel from the numerator and denominator of the travel itinerary choice probabilities, because all travel itineraries have a single non-work stay, and hence these parameters can be set equal to one without loss of generality. In this special case, we recalculate bilateral travel times under the assumption that all consumption travel occurs directly from home. We re-estimate the travel time parameters (ρ_W, ρ_N, θ) using the same moments as for our baseline model above, except for the first set of moments (fraction of multi-stay travel itineraries), for which this special case mechanically predicts no variation.

5.5 Estimation Results

Our baseline sample includes 1,100 locations, comprising 881 Oazas within the 23 wards of central Tokyo and 219 suburban municipalities. We aggregate non-work stays within each location into a single non-work stay for that location. For computational time reasons, we limit the number of stays up to four per day in our baseline model, which covers nearly 90% of observed travel itineraries, as shown in Table 1.

Table 2 reports the estimated parameter values for both our baseline travel itinerary model and the special case with only direct consumption trips. For our baseline model, we find a value for the preference dispersion parameter that determines the elasticity of travel itinerary probabilities with respect to their cost of $\theta = 3.43$. This estimate is comparable to those of the preference dispersion parameter in commuting models, which range from 2.18 to 8.3 in Ahlfeldt *et al.* (2015), Dingel and Tintelnot (2023), Severen (2023) and Kreindler and Miyauchi (2023), although none of these studies incorporates travel itineraries.

We find estimated elasticities of travel cost to travel time of $\rho_W = 0.85$ and $\rho_N = 0.58$. Combined with $\theta = 3.43$, these parameter values imply that a one percent increase in travel time leads to a $2.92 \approx \rho_W \times \theta$ and $1.99 \approx \rho_N \times \theta$ percent decrease in the probability of choosing a travel itinerary on work and non-work days, respectively. This sharp decline in travel with travel time is in line with findings in other contexts, such as Couture (2016), Nevo and Wong

Table 2: Parameter Estimates (Baseline Model and Direct Consumption Trips Model)

Parameter	(1) Baseline	(2) Direct Consumption Trips
σ	4.60	4.60
θ	3.43 (0.08)	3.50 (0.25)
ρ_W	0.85 (0.02)	0.45 (0.03)
ρ_N	0.58 (0.01)	0.45 (0.03)
η_W	1.56 (0.02)	NA
η_N	2.58 (0.05)	NA

Notes: Elasticity of substitution (σ) determined using the ratio of variable profits and revenues in the retail sector; travel parameters estimated using SMM, including workday travel cost (ρ_W); non-workday travel cost (ρ_N); workday intermediate-stay cost (η_W); non-workday intermediate-stay cost (η_N); and dispersion of idiosyncratic preferences (θ); moments weighted using the efficient weighting matrix for the SMM estimates; η_W and η_N omitted from the direct consumption trips model, since they cancel from travel probabilities when all trips have a single non-work stay; standard errors in parentheses computed using the asymptotic variance-covariance matrix.

(2019), and [Davis et al. \(2019\)](#). Our finding that the travel cost parameters are higher during workdays than non-workdays is consistent with the interpretation that the value of time is higher during workdays (e.g., [Buchholz et al. 2025](#)).

We find positive and sizable intermediate-stay costs for each location at which non-traded services are consumed ($\eta_W = 1.56$ and $\eta_N = 2.58$). These sizable intermediate-stay costs are required to rationalize the feature of the data that most agents only visit a few locations during each travel itinerary, despite the love of variety from visiting multiple locations to consume varieties of non-traded services.

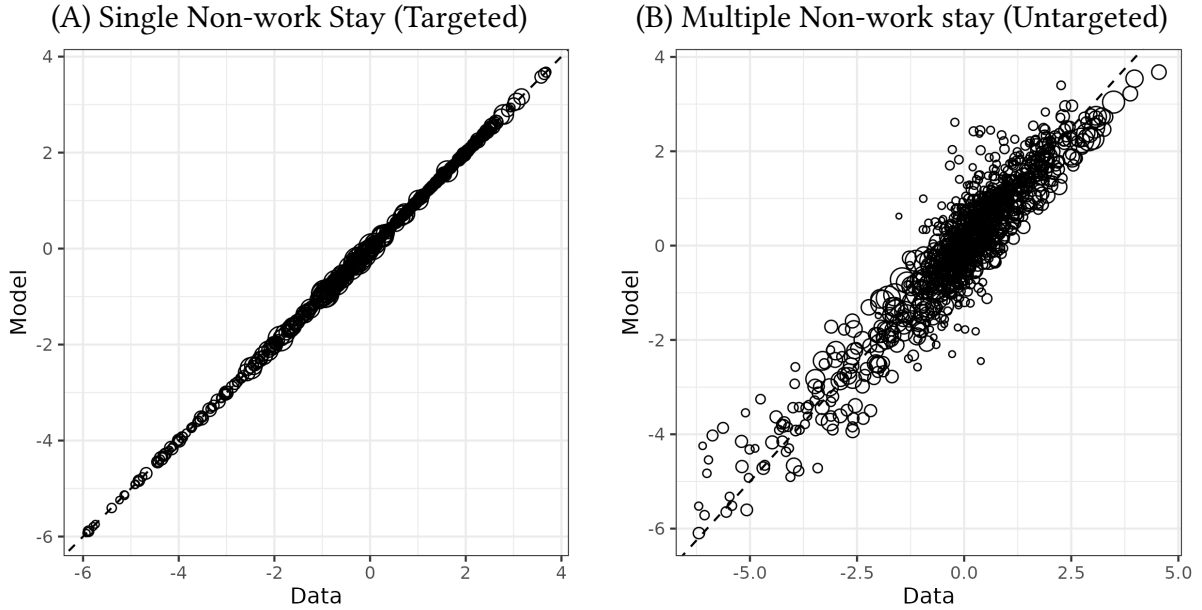
Comparing the parameter estimates for our baseline model and the special case with only direct consumption trips, we find a similar estimated preference dispersion parameter (θ) in both specifications. But we find lower estimated elasticities of travel costs to travel time (ρ_W and ρ_N) with only direct consumption trips, especially on workdays. This pattern reflects the fact that the assumption that all consumption occurs from home leads to longer travel times than in reality, because the consumption of non-traded services often occurs in practice as part of a travel itinerary. To rationalize these longer travel times, the specification with only direct consumption trips requires lower elasticities of travel costs to travel times. Therefore, abstracting from travel itineraries leads to incorrect inference about the sensitivity of travel costs to travel times. In [Section 7.3](#), we show that this underestimation is consequential for the evaluation of transportation infrastructure improvements.

5.6 Model Fit

In Panel (A) of [Figure 4](#), we show the number of travel itineraries with a single non-work stay for each location. The vertical axis shows the value of this moment in the simulated data from the model. The horizontal axis shows the value of this moment in our smartphone

data. Observations correspond to locations, where the size of each circle is proportional to the overall number of non-work stays in a location in our smartphone data. Since we solve for non-traded price indexes in each location to exactly match this moment, all observations lie along the 45-degree line.

Figure 4: Model Fit for Non-Work Stays



Note: Panel A shows the log number of non-work stays per unit area in each location for travel itineraries with a single non-work stay in addition to home and work (targeted); Panel B shows the log number of non-work stays per unit area in each location for travel itineraries with more than one non-work stay in addition to home and work (non-targeted); in both panels, vertical axis shows model predictions, and horizontal axis shows the empirical moment in our smartphone data; each circle corresponds to a location; the size of each circle is proportional to the overall number of non-work stays in a location in our smartphone data. See Appendix Figure F.2.1 for a version of Panel (B) separately for workdays and non-workdays. See Appendix Tables F.2.1 and F.2.2 for the other targeted moments used in our estimation.

In Panel (B) of Figure 4, we provide a specification check on our model's predictions. We show the number of travel itineraries with more than one non-work stay for each location in both the simulated data from the model and our smartphone data. This moment is not targeted in our estimation, because we solve for the non-traded price indexes using the subset of travel itineraries with only one non-work stay. Nevertheless, the model's predictions provide a good approximation to the data, with observations concentrated around the 45-degree line. Naturally, our model makes a number of simplifying assumptions, and hence the model's fit is not expected to be perfect. Nevertheless, we find that we are able to use the structure of the model to successfully extrapolate from the subset of travel itineraries with one non-work stay

to those with more than one non-work stay.¹⁹

5.7 Extended Gravity

Travel itineraries imply that travel decisions are interdependent across the legs of these itineraries. This interdependence leads to a form of extended gravity, in which bilateral travel probabilities depend not only on the bilateral travel time from an origin to a destination, but also on the travel time of the destination from home and work, because destinations are often visited as part of an itinerary including home and work.

We now report an additional specification check, in which we examine the ability of our estimated model to capture this extended gravity property. We begin by establishing that extended gravity holds in our smartphone data, using the observed sequences of non-work stays. We compute the probability that an agent with a given home and work travels from an origin to a destination for a non-work stay. We estimate a Poisson Pseudo Maximum Likelihood (PPML) regression of this travel probability on the log travel times of the destination from the origin, home and work. Observations are combinations of home-work-origin-destination. We use a 10 percent random sample of these observations to keep the sample size manageable. We include fixed effects for destinations and combinations of home-work-origin, in order to control for the multilateral resistance effects emphasized by our travel itinerary model. Identification comes from variation in bilateral travel times for agents with the same home, work, and origin. We cluster the standard errors at the level of the origin location.

Column (1) of Table 3 reports the estimation results using our smartphone data. We find negative and statistically significant coefficients for all three travel times that are of similar magnitudes to one another.²⁰ Column (2) re-estimates the same specification using the simulated data from our estimated model. We find that our estimated model successfully replicates this pattern of negative and statistically significant coefficients of all three travel times. Since our model is an abstraction, and these moments were not targeted in our estimation, it does not perfectly match the exact numerical value of the estimated coefficients, but it clearly captures extended gravity.²¹ Column (3) re-estimates the same specification using the simulated

¹⁹In Online Appendix Figure F.2.1, we show model fit for travel itineraries with a single non-work stay for workdays and non-workdays separately. Although the moment targeted in our estimation pools both sets of observations, we find that the model provides a good fit to the data for these two types of days separately.

²⁰These results are consistent with Davis *et al.* (2019), which uses data on restaurant reviews to show that visits to restaurants are negatively related to distance from home and work. Our smartphone data include all non-work trips, and we observe the sequence in which origins and destinations are visited, which allows us to include home-work-origin and destination fixed effects to control for multilateral resistance effects.

²¹As a further specification check, we estimated an analogous gravity equation using only data on Home-Nonwork-Home travel itineraries. Our theoretical model implies that distance from work should be irrelevant for these travel itineraries, because the non-work destination is visited from home. Consistent with this prediction, we find an estimated coefficient on distance from work that is close to zero and statistically insignificant.

Table 3: Extended Gravity (Untargeted)

Dependent Variable:	Prob Non-Work Trip		
	(1)	(2)	(3)
Model:	Data	Model (Baseline)	Model (Direct Consumption Trips)
log Travel Time (Origin-Dest)	-1.33 (0.02)	-0.67 (0.01)	0.02 (0.01)
log Travel Time (Home-Dest)	-1.82 (0.02)	-1.44 (0.01)	-3.17 (0.00)
log Travel Time (Work-Dest)	-1.81 (0.02)	-1.84 (0.01)	-0.01 (0.01)
Fixed Effects			
Origin-Home-Work	Yes	Yes	Yes
Destination	Yes	Yes	Yes
Observations	43235946	163840600	17411900

Note: For the subset of destinations where non-traded services are consumed, we compute the probability that an agent with a given home and work travels from an origin to a destination. Observations are combinations of home-work-origin-destination during workdays. We use a 10 percent random sample of these observations to keep the sample size manageable. We estimate a Poisson Pseudo Maximum Likelihood (PPML) regression of the probability that an agent travels from an origin to a destination on the log travel times of the destination from the origin, home and work. All specifications include fixed effects for destinations and combinations of home-work-origin. Column (1) uses our smartphone data. Column (2) uses simulated data from our estimated travel itinerary model. Column (3) uses simulated data from the special case of our model in which all consumption travel is assumed to occur from home as direct trips. In this direct consumption trips specification, we set origin-destination travel time equal to work-destination travel time if the non-work stay occurs after work during the day, which ensures that the origin-destination and home-destination travel times are not perfectly collinear. Standard errors are clustered at the level of the origin location.

data from the special case of our model with only direct consumption trips. We find that this special case is unsuccessful in matching extended gravity. Since this special case assumes that all consumption trips occur directly from home, the coefficient on travel time from home is an order of magnitude larger than the coefficients on the other travel times. In contrast, the coefficients on the other travel times are much smaller in magnitude and close to zero. In Section 7, we demonstrate that this failure of the model with only direct consumption trips to capture extended gravity is consequential for counterfactual simulations.

6 Quantitative Urban Model

We now assess the quantitative relevance of the consumption externalities caused by travel itineraries for the organization of economic activity within cities. We embed our specification of travel itineraries in a quantitative general equilibrium urban model that nests conventional commuting models as a special case.²²

We consider a city that consists of a discrete set of locations (city blocks) that are indexed by $h, j \in \{1, \dots, \mathcal{N}\}$. In our baseline specification, we consider a closed city with an exogenous

²²We build on the conventional commuting model of [Ahlfeldt et al. \(2015\)](#). See [Redding and Rossi-Hansberg \(2017\)](#), [Redding \(2023\)](#), and [Redding \(2025\)](#) for reviews of the existing quantitative urban literature.

measure of agents.²³ We normalize this measure of agents to one, such that we solve for the probabilities with which agents make different choices. Each agent is endowed with one unit of labor that is supplied inelastically.

We assume that agents have idiosyncratic preferences for each combination of home, workplace and sector. Each agent observes the realization of these idiosyncratic preferences and picks her preferred combination of home, workplace and sector. She makes this choice of home-workplace-sector before observing the realizations of her idiosyncratic preferences over travel itineraries each day, as modelled in the previous section.

6.1 Preferences and Commuting

Worker preferences are defined over the consumption of a freely-traded final good, non-traded services, and residential floor space. The indirect utility for worker ω with home h , workplace j and sector k depends on her wage (w_{jk}), the price of the freely-traded good (P_{hT}), overall consumption access ($\tilde{\mathbb{A}}_{hj}$) that corresponds to the inverse of the expected cost of travel itineraries to consume non-traded services, the price of residential floor space (Q_h), amenities (B_h), bilateral commuting costs (κ_{hj}) and an idiosyncratic preference draw ($b_{hjk}(\omega)$). We assume that the indirect utility function takes the following Cobb-Douglas form:

$$U_{hjk}(\omega) = \frac{b_{hjk}(\omega) B_h w_{jk} \tilde{\mathbb{A}}_{hj}^{\alpha_S}}{\kappa_{hj} P_{hT}^{\alpha_T} Q_h^{\alpha_H}}, \quad 0 \leq \alpha_T, \alpha_S, \alpha_H \leq 1, \quad (16)$$

where $\alpha_T + \alpha_S + \alpha_H = 1$; we choose the freely-traded final good as the numeraire ($P_{hT} = 1$).

Overall consumption access ($\tilde{\mathbb{A}}_{hj}$) captures the consumption of non-traded services on both workdays and non-workdays. We define overall consumption access as $\tilde{\mathbb{A}}_{hj} = \mathbb{A}_{hj}^\xi \mathbb{A}_{h\emptyset}^{1-\xi}$, where \mathbb{A}_{hj} and $\mathbb{A}_{h\emptyset}$ are measures of consumption access for workdays and non-workdays, respectively, from equation (7). The parameter ξ captures the fraction of expenditure on nontraded services on weekdays. We set this parameter equal to the fraction of workdays during the week, assuming the same expenditure on nontraded services each day.

We assume that commuting costs depend on travel time according to the same functional form as for consumption travel on work days ($\kappa_{hj} = T_{hj}^{\rho_W} \times T_{jh}^{\rho_W}$) in equation (12), but without the intermediate-stay cost.²⁴ We assume that the idiosyncratic preference shock ($b_{hjk}(\omega)$) is drawn from the following independent Fréchet distribution:

$$F_{hj}(b) = e^{-\delta_{hj} b^{-\phi}}, \quad (17)$$

²³It is straightforward to consider an open-city specification, in which total city population is endogenously determined by population mobility with a wider economy that offers a reservation level of utility \bar{U} .

²⁴Recall that travel costs for consumption are measured relative to the travel costs for direct commuting trips in equation (7). Therefore, we introduce the commuting costs separately in equation (16), which allows us to nest conventional commuting models as a special case without consumption of non-traded services ($\alpha_S = 0$).

where ϕ regulates the dispersion of idiosyncratic preferences, and we allow for the location of the preference draws to depend on home and workplace (δ_{hj}), which captures the fact that some pairs of home and workplace can be more attractive than others.

We allow for residential agglomeration forces by assuming that amenities (B_n) depend on residential fundamentals and externalities. Residential fundamentals (\bar{B}_n) capture features of physical geography that make a location a more or less attractive place to live, independent of neighboring economic activity (e.g., green areas). Residential externalities capture the effects of the local density of residents (R_n/K_n), including negative spillovers such as air pollution and crime, and positive externalities through the availability of urban amenities:

$$B_n = \bar{B}_n \left(\frac{R_n}{K_n} \right)^{\gamma_B}, \quad (18)$$

where γ_B governs the strength of residential externalities.²⁵

Using our assumed extreme value distribution for idiosyncratic preferences (17), the probability that an agent chooses to commute from home h to workplace j in sector k is:

$$\Omega_{hjk} = \frac{\delta_{hj} \left(B_h w_{jk} \tilde{A}_{hj}^{\alpha_S} \right)^\phi (\kappa_{hj} Q_h^{\alpha_H})^{-\phi}}{\sum_{h'=1}^{\mathcal{N}} \sum_{j'=1}^{\mathcal{N}} \sum_{k' \in \{T, S\}} \delta_{h'j'} \left(B_{h'} w_{j'k'} \tilde{A}_{h'j'}^{\alpha_S} \right)^\phi (\kappa_{h'j'} Q_{h'}^{\alpha_H})^{-\phi}}, \quad (19)$$

where recall that the traded good is the numeraire ($P_{hT} = 1$ for all h).

Summing across workplaces j and sectors k for each home h , we obtain total residents for home h . Similarly, summing across homes h for each workplace j and sector k , we obtain employment for workplace j and sector k :

$$R_h = \sum_{j=1}^{\mathcal{N}} \sum_{k \in \{T, S\}} \Omega_{hjk}, \quad L_{jk} = \sum_{h=1}^{\mathcal{N}} \Omega_{hjk}, \quad (20)$$

where we used our normalization of a unit measure of agents.

6.2 Production Technology

Both non-traded services and the traded good are produced using labor and floor space.

Nontraded Services We assume that non-traded services are produced under conditions of monopolistic competition and CES demand. We assume a homothetic production technology,

²⁵To focus on travel itineraries as the sole source of spillovers across locations, we assume that residential externalities depend only on a location's own residents' density. However, it is straightforward to incorporate spillovers of residential externalities across locations into equation (18).

such that the fixed and variable costs use the two factors of production with the same factor intensity. The total costs of producing $x_j(\nu)$ units of output of variety ν in location j are:

$$c_j(\nu) = \left(f_{jS} + \frac{x_j(\nu)}{a_{jS}} \right) w_{jS}^{\beta_S} q_j^{1-\beta_S}, \quad 0 < \beta_S < 1,$$

where q_j denotes the price of commercial floor space; f_{jS} parameterizes the fixed cost; a_{jS} is nontraded productivity in location j ; and β_S determines the labor share.

From the first-order condition for profit maximization, the price of non-traded services is a constant markup over marginal costs: $p_j = \frac{\sigma}{\sigma-1} w_{jS}^{\beta_S} q_j^{1-\beta_S} / a_{jS}$. Using profit maximization and zero profits, the measure of nontraded varieties (M_{jS}) produced in location j is:

$$M_{jS} = \frac{1}{\sigma-1} \frac{1}{f_{jS}} \left(\frac{L_{jS}}{\beta_S} \right)^{\beta_S} \left(\frac{H_{jS}}{1-\beta_S} \right)^{1-\beta_S}, \quad (21)$$

where L_{jS} and H_{jS} are the aggregate inputs of labor and commercial floor space used in nontraded services in location j .

Combining expenditure minimization, profit maximization and zero profits, we obtain the following solution for the non-traded services price index (P_{jS}) that entered into travel itinerary decisions in the previous section:

$$P_{jS} = p_j (M_{jS})^{\frac{1}{1-\sigma}} = \frac{1}{A_{jS}} w_{jS}^{\beta_S} q_j^{1-\beta_S}, \quad (22)$$

where A_{jS} is a supply shifter for each location that depends on productivity and aggregate factor inputs of labor and commercial floor space in the nontraded services sector:²⁶

$$A_{jS} \equiv \tilde{a}_{jS} (L_{jS})^{\frac{\beta_S}{\sigma-1}} (H_{jS})^{\frac{1-\beta_S}{\sigma-1}}, \quad (23)$$

$$\tilde{a}_{jS} \equiv \frac{\sigma-1}{\sigma} \left[\frac{1}{\sigma-1} \left(\frac{1}{\beta_S} \right)^{\beta_S} \left(\frac{1}{1-\beta_S} \right)^{1-\beta_S} \right]^{\frac{1}{\sigma-1}} \left(\frac{1}{f_{jS}} \right)^{\frac{1}{\sigma-1}} a_{jS}.$$

Traded Good The traded good is produced under conditions of perfect competition using labor and floor space. We assume a constant returns to scale Cobb-Douglas production technology. Profit maximization and zero profits imply that the price of the traded good in location j must equal its unit cost in all locations with positive production:

$$P_{jT} = \frac{1}{A_{jT}} w_{jT}^{\beta_T} q_j^{1-\beta_T}, \quad 0 < \beta_T < 1, \quad (24)$$

²⁶Although we assume monopolistically-competitive nontraded service firms, this formulation is isomorphic to an alternative specification of perfectly competitive nontraded service firms with agglomeration forces from external economies of scale, analogous to equation (25) below.

where A_{jT} denotes traded sector productivity and recall that $P_{jT} = 1$.

We incorporate agglomeration forces in the traded sector by allowing productivity (A_{jT}) to depend on production fundamentals and production externalities. Production fundamentals (\bar{A}_{jT}) capture features of physical geography that make a location more or less productive independently of neighboring economic activity (e.g., access to natural water). Production externalities capture productivity benefits from the local density of employment. Formally, we assume that productivity for the traded good satisfies:

$$A_{jT} = \bar{A}_{jT} \left(\frac{L_{jT}}{K_j} \right)^{\gamma_A}, \quad (25)$$

where K_j denotes area and γ_A governs the strength of production externalities.²⁷

6.3 General Equilibrium

We assume that demand equals supply in each location in the markets for residential floor space, commercial floor space, commuters, and non-traded goods, as detailed in Appendix G.1. The general equilibrium of the model is referenced by the inverse of the expected cost of travel itineraries to consume non-traded services (\tilde{A}_{hj}); the travel itinerary choice probabilities ($\Lambda_{I|hj}$); the residence-workplace-sector choice probabilities (Ω_{hjk}); the price index for nontraded services (P_{jS}); the prices for commercial and residential floor space (q_j, Q_j); and the wage (w_{jk}). The equilibrium values of these variables are determined by consumers' optimal itinerary decisions in equations (6) and (7); workers' choices of workplace-residence-sector from equations (19) and (20); firms' optimal decisions in each sector from equations (22) and (24); the amenity and productivity spillovers (18) and (25); and the market clearing conditions.

To undertake counterfactuals, we follow an exact-hat algebra approach (Dekle *et al.* 2007), which uses the initial values of the endogenous variables and model parameters. Given an assumed change in the exogenous variables of the model (e.g., travel times, τ_{in}), the counterfactual equilibrium can be computed using the values of the structural parameters $\{\sigma, \xi, \theta, \rho_W, \rho_N, \eta_W, \eta_N, \phi, \alpha_T, \alpha_S, \alpha_H, \beta_S, \beta_T, \gamma_B, \gamma_A\}$ and the initial values of the endogenous variables of the model $\{P_{jS}, w_{jk}, \Omega_{hjk}, R_h, L_{jk}\}$. In Online Appendix G.2, we report the full system of equations for the counterfactual equilibrium.

6.4 Parameterization

We use the travel itinerary parameters $\{\sigma, \theta, \rho_W, \rho_N, \eta_W, \eta_N\}$ and nontraded price indexes $\{P_{nS}\}$ from the estimation of our travel itinerary model in Section 5.4 above.

²⁷Again to focus on travel itineraries as the sole source of spillovers across location, we assume production externalities depend only on a location's own employment density. But it is straightforward to incorporate spillovers of production externalities across locations into equation (25).

We calibrate the remaining general equilibrium parameters $\{\alpha_T, \alpha_S, \alpha_H, \phi, \xi, \beta_S, \beta_T, \gamma_B, \gamma_A\}$ using our Japanese data and central values from the existing empirical literature, as summarized in Table 4 below. We calibrate the Cobb-Douglas expenditure share parameters $\{\alpha_T, \alpha_S, \alpha_H\}$ using aggregate data on expenditure shares in Japan. We set the share of expenditure on residential floor space equal to $\alpha_H = 0.25$ from [Statistics Bureau of Japan \(2020\)](#). We set the expenditure share for nontraded services as $\alpha_S = 0.47$, based on the revenue share of those sectors in [Statistics Bureau of Japan \(2016\)](#).²⁸ We recover the implied share of expenditure on traded goods as $\alpha_T = 1 - \alpha_H - \alpha_S$.

Table 4: Calibration of General Equilibrium Parameters

Parameter	Value	Description	Source
α_H	0.25	Expenditure share for housing	Household Expenditure Survey
α_S	0.47	Expenditure share for nontradable service	Household Expenditure Survey
β_S	0.81	Share of labor in nontradable service sector	Economic Census
β_T	0.82	Share of labor in tradable sector	Economic Census
ϕ	3.43	Dispersion of preference shocks	Set to $\phi = \theta$
γ_A	0.28	Agglomeration productivity spillover elasticity	Set to $1/(\sigma - 1)$
γ_B	0.00	Agglomeration amenity spillover elasticity	Range of Values
ξ	0.71	Fraction of workdays (pre-COVID-19)	Set to 5/7

Note: The list of general equilibrium parameters and their calibrated values. We use the estimated travel itinerary parameters for our baseline travel itinerary model from Table 2.

We calibrate the shares of commercial floor space in costs in non-traded services ($1 - \beta_S$) and traded goods ($1 - \beta_T$) as equal to 0.19 and 0.18, respectively, from [Statistics Bureau of Japan \(2016\)](#). We set the dispersion of idiosyncratic preferences for commuting equal to the value estimated from our travel itinerary model ($\phi = \theta$). We set the agglomeration forces parameter for the traded sector equal to the value implied by the love of variety and increasing returns to scale in non-traded services, such that both sectors have the same agglomeration forces ($\gamma_A = 1/(\sigma - 1)$). We abstract from residential externalities ($\gamma_B = 0$) in our baseline specification to focus on travel itineraries as the sole source of consumption externalities. We report the sensitivity of our results to alternative parameter values, based on the range of estimates reported in the meta-analyses of estimated agglomeration forces in [Melo et al. \(2009\)](#) and [Ahlfeldt and Pietrostefani \(2019\)](#). We calibrate the fraction of workdays (ξ) at 5/7 before the COVID-19 pandemic, based on the regular working schedule in Japan and the fact that working from home was rare in Japan before the pandemic ([Okubo 2022](#)).

We obtain the values of commuting probabilities, residents and employment in the initial equilibrium $\{\Omega_{hjk}, R_h, L_{jk}\}$ by multiplying the commuting flows from our smartphone data by

²⁸The nontraded service sector is defined as including the following Japanese Standard Industrial Classification (JSIC) categories: “Finance, Real Estate, Communication, and Professional” (G, J, K, L); “Wholesale and Retail” (I), “Accommodations, Eating, Drinking” (M), “Medical and Health Care” (P), and “Other Services” (Q).

the sectoral employment share at each employment location from the economic census. This approach is consistent with the prediction of our model that commuting probabilities do not vary by sector conditional on workplace.²⁹ We obtain wages in the initial equilibrium $\{w_{nk}\}$ from the model’s market clearing conditions, as discussed further in Online Appendix G.3.

7 Counterfactuals

We now use our quantitative general equilibrium model to undertake counterfactuals to show that travel itineraries are key to explaining the decline of economic activity in central areas following the shift to WFH, driving the agglomeration of economic activity in these areas, and providing a rationale for infrastructure policies in these areas.

7.1 Working from Home (WFH)

We first provide further validation of our estimated travel itinerary model by comparing its predictions for an exogenous shock outside our estimation sample to the observed impact of this shock in the data. We undertake a counterfactual for WFH and show that our general equilibrium urban model with travel itineraries is quantitatively successful in matching our quasi-experimental evidence for changes in travel patterns following the shift to WFH. In contrast, the special case of a conventional urban model in which all consumption occurs through direct trips from home is unsuccessful in matching these empirical findings.³⁰

We interpret the shift to WFH in our model of travel itineraries as a change in two sets of parameters. First, WFH reduced the number of days each week that people commute into work (ξ). We set this reduction in the fraction of commuting days (ξ) at 20 percent, consistent with the survey data for Tokyo (Okubo 2022; Online Appendix B.2). We assume that WFH does not directly reduce the share of commercial floor space in production costs in each sector ($1 - \beta_S, 1 - \beta_T$), because under hybrid working commercial floor space is required on the days on which workers commute into work each week. Second, several commentators have argued that the shift to WFH changed workers’ preferences for the characteristics of residential locations, such as higher demand for larger residences and suburban natural amenities. This change in preferences can be captured by a change in the relative amenities of locations (\bar{B}_n). Therefore, we allow residential amenities in each location (\bar{B}_n) to change between our baseline sample in 2019 and 2023, such that each model exactly matches the observed change in the

²⁹We set $\Omega_{hjk} = (\sum_{k'} \Omega_{hjk'}) \times \frac{L_{jk}}{\sum_{k'} L_{jk'}}$, where we obtain the first term from our smartphone data and the second term from the economic census. From equation (19), our model implies this second term does not depend on home h , because commuting flows only depend on sector through workplace.

³⁰We focus on our quasi-experimental findings for WFH, because our event studies for retail store closure in Section 4.3 are estimated at a finer level of spatial aggregation than our travel itinerary model.

number of residents in each location between these years. Although both our model and the conventional urban model exactly match observed residents, they need not match observed patterns of work and non-work trips.

Given these assumed changes in parameters $(\hat{\xi}, \hat{B}_n)$, we solve for a counterfactual equilibrium in our baseline travel itinerary model and the conventional urban model, holding all other parameters constant.³¹ In Panel (A) of Table 5, we report the mean change in non-work stays for the same four groupings of locations as used for our quasi-experimental findings in Figure 2: CBD, high-density, medium-density and low-density. We report these mean changes in our smartphone data, our baseline travel itinerary model, and the special case of a conventional urban model with only direct consumption trips. We find that our baseline model of travel itineraries is quantitatively successful in capturing the larger decline in non-work stays in central locations in the data. As workers commute less frequently into central locations, this reduces non-work trips in those locations, as fewer workers stop off to consume non-traded services (e.g., coffee shops and bars) along the way. In contrast, the conventional urban model with only direct consumption trips is unsuccessful in capturing this pattern in the data. Changes in non-work stays are close to zero in this special case, because all consumption trips are assumed to occur directly from home, regardless of where work takes place.³²

To provide further evidence on the model’s general equilibrium predictions, Panel (B) reports the mean log change in the price of commercial floor space for the same four groups of locations. In the data, we measure these log changes using the office/commercial building rent data from Sanko Estate Company Ltd. (see Online Appendices A.9 and B). In the model, we measure these log changes using the counterfactual predictions of the model for the change in the price of commercial floor space, assuming segmented markets for commercial and residential floor space. Again we report results for both our baseline model and the special case of a conventional urban model. To abstract from inflation in the observed data and to focus on relative changes, we normalize the mean log change in the price of commercial floor space by the value for low-density locations, such that low-density locations have a value of zero.

We find that our baseline model of travel itineraries is quantitatively successful in capturing the larger decline in commercial floor space prices in central locations. As workers commute less frequently into central locations, and fewer workers stop to consume non-traded

³¹In principle, the shift to WFH could have changed the travel parameters $\{\rho_W, \rho_N, \eta_W, \eta_N\}$. If we re-estimate our travel itinerary model using 2023 data, we find similar parameter values as for our baseline 2019 sample (Online Appendix Table H.1.3). This pattern of results is consistent with the idea that WFH mainly changed the number of days each week people commute to work rather than travel behavior conditional on commuting.

³²Changes in non-work stays in the special case with only direct consumption trips are not exactly zero, because the travel elasticities are different between workdays and non-workdays ($\rho_W \neq \rho_N$), and we shift the fraction of workdays (ξ). However, our estimates of these travel elasticities are similar to one another (Column (2) of Table 2), and hence this special case predicts little change in non-work stays.

Table 5: Actual and Counterfactual Changes in Non-work Stays, Commercial Floor Space Prices and the Number of Non-traded Varieties Following the Shift to WFH

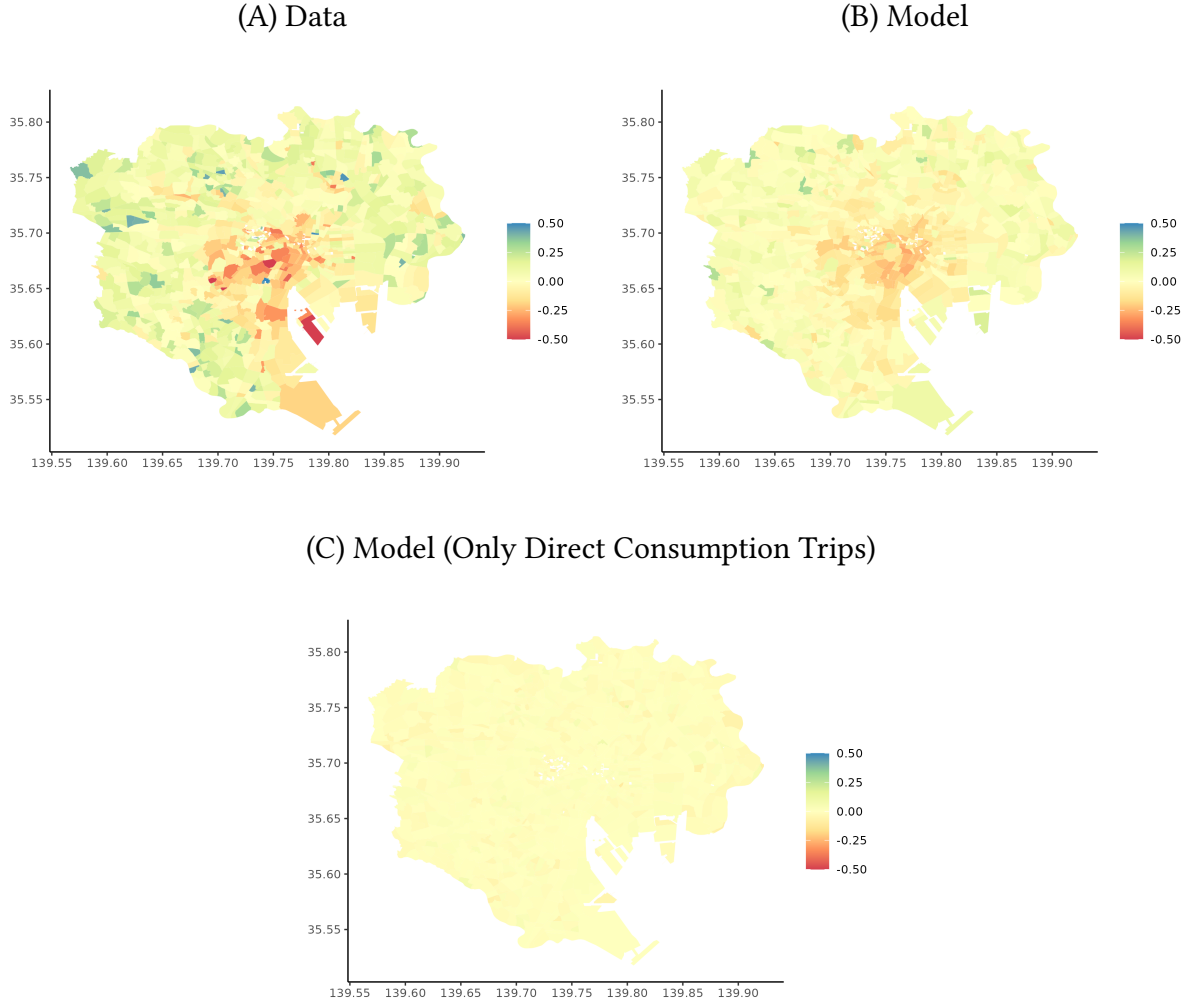
	CBD	High	Medium	Low
(A) Non-Work Stays				
(1) Data	-0.18	-0.11	0.04	0.03
(2) Model: Baseline	-0.15	-0.11	0.02	0.04
(3) Model: Only direct consumption trips	0.00	0.00	0.00	0.00
(B) Commercial floor space price (normalized)				
(1) Data	-0.12	-0.11	-0.02	[0.00]
(2) Model: Baseline	-0.13	-0.09	-0.01	[0.00]
(3) Model: Only direct consumption trips	0.00	0.00	0.00	[0.00]
(C) Nontraded service varieties (normalized)				
(1) Data	-0.15	-0.09	-0.05	[0.00]
(2) Model: Baseline	-0.11	-0.08	-0.02	[0.00]
(3) Model: Only direct consumption trips	0.00	0.00	0.00	[0.00]

Note: “(1) Data” corresponds to the observed patterns in our data (see the main text and Online Appendix A.9 and B for the data sources); “(2) Model: Baseline” is our estimated model of travel itineraries; “(3) Model: Only direct consumption trips” is the special case of a conventional urban model, in which all consumption trips occur directly from home; Panel (A) reports mean changes in non-work stays; Panel (B) report mean log changes in commercial floor space prices; Panel (C) reports mean log changes in the number of non-traded varieties; in Panels (B) and (C), commercial floor space prices and the number of non-traded varieties normalized such that the mean for low-density locations is zero (as indicated by [0.00]); central business district (CBD) corresponds to within 2km of the centroid of the Oaza of Chiyoda ward; high-density corresponds to locations with employment densities in the top 10 percent (excluding the CBD); mid-density corresponds to locations with employment densities from the 50-90th percentiles (excluding the CBD); low-density corresponds to locations with employment densities from the 0-50th percentiles (excluding the CBD).

services along the way, this reduces the demand for commercial floor space, and leads to a decline in commercial floor space prices. The lower demand for non-traded services in the central city also leads to a reduction in the variety of these services, thereby reducing the attractiveness of commuting into the central city, and further depressing the demand for non-traded services. In contrast, the conventional urban model with only direct consumption trips implies little change in the demand for non-traded services, because all consumption occurs directly from home, regardless of where work takes place.

In Panel (C), we provide further evidence on these general equilibrium predictions by reporting the mean log change in the number of non-traded service varieties. In the data, we measure these log changes using the number of service establishments in the phone directory data. In the model, we measure these log changes using the counterfactual changes in the number of non-traded service varieties. Again, to abstract from macro shocks in the observed data and to focus on relative changes, we normalize the mean log changes by the value for low-density locations. In the data, we find larger declines in non-traded varieties in central

Figure 5: Actual and Counterfactual Declines in Non-work Stays Following the Shift to WFH



Note: Log Change in non-work stays following the Shift to WFH in Oazas within 23 central wards (municipalities) in Tokyo Metropolitan Area; Panel (A) shows log changes in the data from 2019-2023; Panel (B) shows the counterfactual predictions of our estimated model of travel itineraries; Panel (C) shows the counterfactual predictions of the special case of the conventional urban model without travel itineraries.

locations. Our estimated model of travel itineraries is quantitatively successful in capturing this pattern. Again, we find that the special case of the conventional urban model with only direct consumption trips is unsuccessful at capturing the observed changes in the data.

In Figure 5, we provide greater spatial resolution on the changes in non-work stays in the 23 wards (municipalities) of Central Tokyo following the shift to WFH. We show these changes in non-work stays in our smartphone data (Panel (A)); our estimated model of travel itineraries (Panel (B)); and the special case of a conventional urban model without travel itineraries (Panel (C)). We find that our model of travel itineraries is relatively successful at capturing the spatial pattern of the changes in non-work stays in the data. In both the model and data, we find

that the largest declines in non-work stays are concentrated in the center of Tokyo’s business district, although the absolute magnitude of the changes is somewhat smaller in the model. In contrast, the conventional urban model generates changes in non-work stays that are close to zero throughout the metropolitan area, because it assumes that all consumption occurs through direct trips from home, regardless of where work takes place.

In Online Appendix [H.1](#), we provide further sensitivity analyses of these results. We show that we obtain similar counterfactual predictions across the following specifications: (i) re-estimating the travel itinerary parameters $\{\theta, \rho_W, \rho_N, \eta_W, \eta_N\}$ using our smartphone data for 2023; (ii) calibrating our model using smoothed commuting flows in the initial equilibrium ([Dingel and Tintelnot 2023](#)); and (iii) using alternative assumed parameters for production agglomeration forces (γ_A). We also demonstrate a similar pattern of results in a restricted specification assuming only one non-work stay in addition to work per travel itinerary (instead of up to four non-work stays in our baseline specification). Therefore, to explain the observed impact of the shift to WFH in the data, it is critical to model the complementarity between work and non-work stays through travel itineraries. In contrast, the complementarity between multiple non-work stays plays a more modest role for the shift to WFH. As a final specification check, we show that our travel itinerary model is also successful in capturing the decline in employment in the CBD following the shift to WFH.

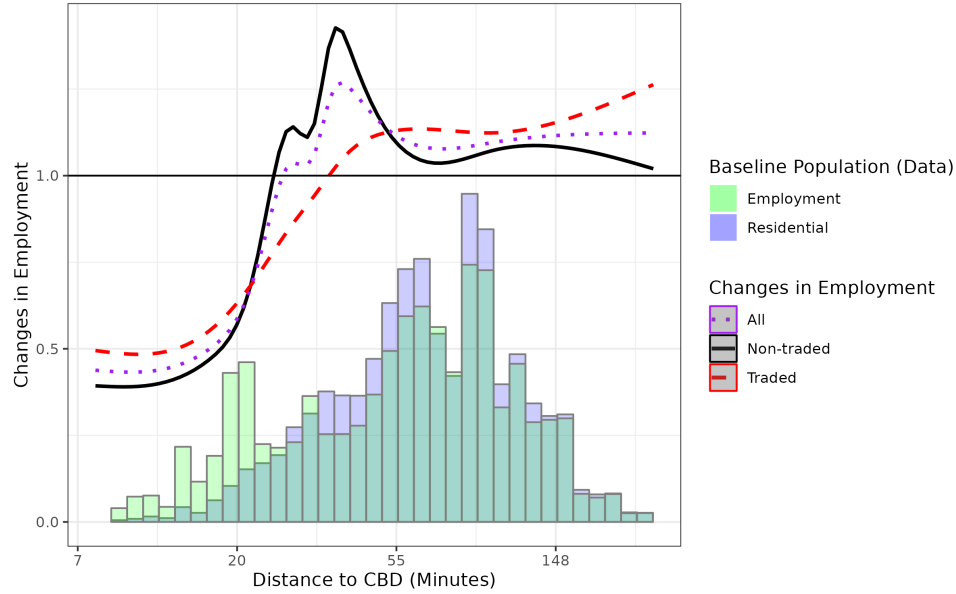
7.2 Travel Itineraries and Agglomeration

We next use our estimated travel itinerary model to show that the consumption externalities from travel itineraries are a quantitatively important force for the agglomeration of economic activity in central cities. Starting from the observed equilibrium in the data in 2019, we undertake a counterfactual in which we assume that intermediate-stay costs become prohibitive, and instead assume that agents make one direct trip from home to consume non-traded services and another direct trip to commute into work.

Figure [6](#) displays the results from this counterfactual. The three lines show the fitted values from local polynomial regressions of counterfactual log changes in employment on log travel time to the CBD. The solid, dashed and dotted lines represent values for the non-traded services sector, the traded sector, and total employment, respectively. As a point of comparison, we overlay on the same figure histograms of the distribution of employment (light green) and residents (blue) in the initial equilibrium by distance bin from the CBD.

We find that the consumption externalities from multi-stay travel itineraries (trip chains) play a quantitatively important role in the concentration of employment in Central Tokyo. In the most central parts of the city (within 20 minutes from the CBD), overall employment by workplace falls by up to around one half. We find reductions in employment of a similar

Figure 6: Counterfactual Employment Changes Without Multi-Stay Travel Itineraries



Note: Counterfactual for prohibitive intermediate-stay costs, such that agents instead make one direct trip from home to consume non-traded services, and another direct trip to commute into work (no trip chains); three lines show fitted values from local polynomial regressions of counterfactual log changes in employment on log travel time to the CBD; solid, dashed and dotted lines represent values for the non-traded services sector, the traded sector, and total employment, respectively; histograms show the distribution of employment (light green) and residents (blue) in the initial equilibrium by distance bin from the CBD.

magnitude for both the non-traded and traded sectors. This fall in employment in central areas is compensated by an increase in nontraded sector employment in the inner suburbs (within 55 minutes from the CBD) and an increase in traded sector employment (and to a lesser extent, non-traded sector employment) in the city's outskirts (outside 55 minutes from the CBD).

This pattern of results is intuitive. The fall in employment in the nontraded sector in central areas occurs because of the loss in demand for non-traded services from in-commuters through their travel itineraries. The fall in employment in the traded sector in central areas occurs because the attractiveness of these areas as workplaces relies substantially on the presence of nontraded services offered along commuters' travel itineraries. These two forces together generate a positive feedback loop between the nontraded and traded sectors, thereby providing a force for agglomeration in the central city.

In Online Appendix [H.2](#), we undertake additional counterfactual simulations, in which we restrict travel itineraries to include only one non-work stay in addition to work. We find that the counterfactual changes in employment exhibit a similar pattern but are smaller in magnitude. This pattern of results is consistent with the interpretation that a key agglomeration force is created by the complementarity between work and non-work stays within the same travel itinerary. The complementarity between multiple non-work stays also plays a role, but

is smaller in magnitude, in part because of the substitution of spending across multiple non-work stays within the same travel itinerary.

These results provide an explanation for why cities like Tokyo, London, New York City and Paris have a flourishing city center with a wide range of non-traded services (Brueckner *et al.* 1999, Glaeser *et al.* 2001). Each of these cities has dense networks of public transportation radiating outwards from the central city, which supports large flows of in-commuters. These in-commuters generate demand for non-traded services along their travel itineraries, and the provision of these non-traded services in turn makes the central city more attractive to in-commuters through a mutually-reinforcing process of agglomeration.

7.3 Place-Based Infrastructure Policies

Finally, we show that the consumption externalities from travel itineraries are quantitatively relevant for the evaluation of alternative place-based infrastructure policies, and favor investments in central cities where travel itineraries are more prevalent.

Table 6 reports counterfactuals for two specific transport improvements: (i) the construction of the Yamanote Railway Line, which follows a circular route connecting subcenters of central Tokyo; (ii) the construction of the Chuo Railway Line, which follows a radial route that connects an outer suburb of Tokyo to the center of the city (see Online Appendix Figure H.3.1 for the map). First, we recompute bilateral travel times between locations within Tokyo without each of these railway lines. Second, starting from the observed equilibrium in the data in our baseline sample in 2019, we solve for an exact-hat algebra counterfactual equilibrium without each railway line. Third, we compute the percentage welfare gain from constructing each of these railway lines, which equals welfare in the actual equilibrium with these lines divided by welfare in the counterfactual equilibrium without each line. We abstract from construction costs, because they take the same value across all of the different model specifications that we consider.

We report the results of these counterfactuals for three different model specifications. In row (1), we consider our baseline estimated travel itinerary model. In row (2), we examine a conventional urban model with only direct consumption trips (as in the previous subsection), using the parameter estimates $\{\theta, \rho_W, \rho_N\}$ under this model specification (Table 2; Column 2). In row (3), we consider this conventional urban model with only direct consumption trips, but using the parameter estimates $\{\theta, \rho_W, \rho_N\}$ from our baseline model (Table 2; Column 1). We consider two alternative versions of the model with only direct consumption trips to clarify how the changes in the estimated parameters and the undercounting of the volume of trips in this specification affect the evaluation of alternative transport infrastructures.

Across all three models, we find a modest welfare improvement from the construction

Table 6: Welfare Gains from Alternative Transport Improvements

(A) Yamanote Line (Circular Lines within the Central City)		
Specification	Welfare Gains (%)	Relative to Baseline (%)
(1) Baseline	0.44	100
(2) Only Direct Consumption Trips	0.28	61
(3) Only Direct Consumption Trips (Baseline θ, ρ_W, ρ_S)	0.36	80
(B) Chuo Line (Radial Lines between the Central City and Suburbs)		
Specification	Welfare Gains (%)	Relative to Baseline (%)
(1) Baseline	0.71	100
(2) Only Direct Consumption Trips	0.49	68
(3) Only Direct Consumption Trips (Baseline θ, ρ_W, ρ_S)	0.66	93

Note: Welfare gains from the construction of the Chuo Railway Line (radial route between a suburb and the center) and Yamanote Railway Line (circular route connecting subcenters of central Tokyo); table reports welfare in the actual equilibrium (with these railway lines) relative to welfare in the counterfactual equilibrium (without each railway line); row (1) corresponds to our baseline estimated travel itinerary model; row (2) corresponds to a conventional urban model, in which all travel (commuting and consumption) occurs through direct trips from home, using the parameter estimates $\{\theta, \rho_W, \rho_N\}$ from this model specification (Table 2; Column 2); row (3) corresponds to a conventional urban model with only direct trips, but using the parameter estimates $\{\theta, \rho_W, \rho_N\}$ from our baseline model (Table 2; Column 1); second column reports percentage welfare gains from each transport improvement; third column reports the welfare gain from each transport improvement relative to that our baseline travel itinerary model (as a percentage).

of each railway line, which is consistent with each of these lines involving a relatively small improvement in the overall railway network for Tokyo. For both railway lines, we find that the welfare gain in the conventional urban model in which all consumption trips occur directly from home (row (2)) is only around two-thirds (61-68 percent) of the welfare gains in our baseline estimated travel itinerary model (row (1)). This pattern occurs for two main reasons. First, the conventional urban model undercounts trips that occur within a travel itinerary. Second, the conventional urban model underestimates the elasticity of travel cost to travel time $\{\rho_W, \rho_N\}$ (Table 2). Therefore, for the same reduction in travel time, these models tend to predict lower welfare gains (Small and Verhoef 2007, Donald *et al.* 2024).

To further understand these two margins, in row (3), we examine the conventional urban model with only direct consumption trips using the same estimated travel cost parameters as in our baseline travel itinerary model (row 1). We find that this specification yields higher welfare gains than row (2), suggesting that the downward-biased estimates of $\{\rho_W, \rho_N\}$ contribute to the underestimation of welfare gains using the conventional urban model. However, this specification still underpredicts the welfare gains relative to our baseline model (row 1), as it still undercounts trips that occur within a travel itinerary.

Notably, the extent to which welfare is lower in the conventional urban model differs between the Yamanote Line (80 percent) and the Chuo Line (93 percent). This pattern of results is consistent with the idea that the Yamanote Line is likely to be used much more frequently for travel itineraries than the Chuo Line, because it connects the subcenters of central Tokyo with a high density of non-traded services. In contrast, the Chuo Line runs from the outskirts

of Tokyo to the center, and is less suited for travel itineraries involving multiple locations, because of the lower density of non-traded services in the outer areas of Tokyo.

In Online Appendix H.3, we report additional sensitivity and robustness analyses. Restricting the number of non-work stays within travel itineraries to one (from four non-work stays in our baseline specification) leads to a reduction in the predicted welfare gains, particularly for Yamanote Line, which is used extensively for multi-stay travel itineraries. We find that our results are mostly unchanged if we calibrate our model using smoothed commuting flows or assume alternative values for spillovers in production (γ_A) or amenities (γ_B).

Overall, we find that the consumption externalities implied by travel itineraries are quantitatively relevant for evaluating alternative place-based infrastructure projects, and favor investments in central areas where multi-stay travel itineraries disproportionately occur.

8 Conclusions

We develop a tractable quantitative framework for modelling the rich patterns of spatial mobility observed in smartphone Geographical Positioning System (GPS) data. These data allow us to quantify one of the essential characteristics of cities, namely the millions of travel journeys each day between locations of residence, work and consumption.

We show that commuting and consumption travel are frequently undertaken as part of a travel itinerary, defined as a journey starting and ending at home on a given day that can include more than one stop along the way. These travel itineraries give rise to consumption externalities and both complementarity and substitutability between locations. Although the empirical relevance of these travel itineraries is known, modelling them theoretically is challenging, because of the high-dimensionality of the state space in choosing combinations of locations and sequences in which to visit them.

We begin by providing evidence of consumption externalities using two different sources of quasi-experimental variation. First, we examine the shift to WFH following the COVID-19 pandemic. Second, we use event-study specifications for the closure of large retail stores (more than 5,000 square meters floor space). In both cases, we find strong evidence of consumption externalities, such that as one location becomes less attractive, the resulting reduction in foot traffic spills over to other locations that are nearby or along the way.

We next develop a tractable quantitative model of travel itineraries to rationalize these empirical findings. Agents choose where to live, where to work, and where to consume non-traded services. Agents have love of variety preferences over these non-traded services, but face travel costs that are increasing in travel time, and incur an intermediate-stay cost for each intermediate location at which they consume non-traded services. Travel costs provide the in-

centive for agents to chain trips along travel itineraries. We overcome the high-dimensionality of the state space using importance sampling.

We show that our model is quantitatively successful in explaining the observed changes in economic activity following the shift to WFH in the aftermath of the COVID-19 pandemic. We show that the reduction in work trips to the central city leads to a collapse in non-work trips there, as workers no longer stop off to consume non-traded services along the way to and from work. This reduction in work and non-work trips leads in general equilibrium to a larger decline in the relative price of commercial floor space and the number of non-traded service varieties in central locations. In contrast, a conventional urban model is unsuccessful in capturing these patterns in the data, because workers are assumed to travel to consume non-traded services from home, regardless of where they work.

More broadly, we show that the consumption externalities implied by travel itineraries are central to the agglomeration of economic activity in central cities and the evaluation of alternative place-based infrastructure policies in these central areas.

References

- ACKERBERG, D. A. (2009). A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation. *Quantitative Marketing and Economics*, 7 (4), 343–376.
- AGARWAL, S., BRADFORD JENSEN, J. and MONTE, F. (2020). Consumer Mobility and the Local Structure of Consumption Industries. *NBER Working Paper*, 23616.
- AHLFELDT, G. and PIETROSTEFANI, E. (2019). The economic effects of density: A synthesis. *Journal of Urban Economics*, 111, 97–107.
- , REDDING, S., STURM, D. and WOLF, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83 (6), 2127–2189.
- ALLEN, T. and ARKOLAKIS, C. (2023). The welfare effects of transportation infrastructure improvements. *Review of Economic Studies*, 89 (6), 2911–2957.
- , — and LI, X. (2017). Optimal city structure. Yale University, mimeograph.
- , ATKIN, D., CANTILLO, S. and HERNANDEZ, C. (2021a). Trucks. MIT, mimeograph.
- , FUCHS, S., GANAPATI, S., GRAZIANO, A., MADERA, R. and MONTORIOL-GARRIGA, J. (2021b). Urban welfare: Tourism in barcelona. Dartmouth College, mimeograph.
- ALMAGRO, M. and DOMÍNGUEZ-IINO, T. (2025). Location sorting and endogenous amenities: Evidence from amsterdam. *Econometrica*, 93 (3), 1031–1071.
- ALONSO, W. (1964). *Location and Land Use*. Cambridge MA: Harvard.
- ALTHOFF, L., ECKERT, F., GANAPATI, S. and WALSH, C. (2022). The geography of remote work. *Regional Science and Urban Economics*, 93, 103770.

- ALVAREZ, F. and LUCAS, R. E. (2007). General equilibrium analysis of the eaton–kortum model of international trade. *54* (6), 1726–1768.
- ANTRÀS, P., FORT, T. C. and TINTELNOT, F. (2014). The Margins of Global Sourcing: Theory and Evidence from U.S. Firms. *American Economic Review*, **107** (9), 2514–64.
- ARKOLAKIS, C., ECKERT, F. and SHI, R. (2023). Combinatorial Discrete Choice. *NBER Working Paper*, **31877**.
- ATHEY, S., FERGUSON, B., GENTZKOW, M. and SCHMIDT, T. (2021). Estimating experienced racial segregation in us cities using large-scale gps data. *Proceedings of the National Academy of Sciences*, **118** (46), e2026160118.
- ATKIN, D., CHEN, K. and POPOV, A. (2022). The returns to face-to-face interactions: Knowledge spillovers in silicon valley. *NBER Working Paper*, **30147**.
- BALBONI, C., BRYAN, G., MORTEN, M. and SIDDIQI, B. (2021). Transportation, gentrification, and urban mobility: The inequality effects of place-based policies. MIT, mimeograph.
- BARRERO, J. M., BLOOM, N. and DAVIS, S. J. (2023). The evolution of work from home. *Journal of Economic Perspectives*, **37** (4), 23–50.
- BENMELECH, E., BERGMAN, N., MILANEZ, A. and MUKHARLYAMOV, V. (2019). The Agglomeration of Bankruptcy. *Review of Financial Studies*, **32** (7), 2541–2586.
- BERRY, S., KHWAJA, A., KUMAR, V., MUSALEM, A., WILBUR, K. C., ALLENBY, G., ANAND, B., CHINTAGUNTA, P., HANEMANN, W. M., JEZIORSKI, P. and MELE, A. (2014). Structural models of complementary choices. *Marketing Letters*, **25** (3), 245–256.
- BORUSYAK, K., JARAVEL, X. and SPIESS, J. (2024). Revisiting event study designs: Robust and efficient estimation. *Review of Economic Studies*, **91** (6), 3253–3285.
- BRANCACCIO, G., KALOUPSIDI, M. and PAPAGEORGIOU, T. (2020). Geography, transportation and endogenous trade costs. *Econometrica*, **88** (2), 657–691.
- BRUECKNER, J. K., THISSE, J.-F. and ZENOU, Y. (1999). Why is central paris rich and downtown detroit poor?: An amenity-based theory. *European economic review*, **43** (1), 91–107.
- BÜCHEL, K., EHRLICH, M. V., PUGA, D. and VILADECANS, E. (2020). Calling from the outside: The role of networks in residential mobility. *Journal of Urban Economics*, **119**, 103277.
- BUCHHOLZ, N., DOVAL, L., KASTL, J., MATEJKA, F. and SALZ, T. (2025). Personalized pricing and the value of time: Evidence from auctioned cab rides. *Econometrica*, **93** (3), 929–958.
- CATS (1956). *Chicago Area Transportation Study 1956*. Chicago: State of Illinois.
- CLAYCOMBE, R. J. (1991). Spatial retail markets. *International Journal of Industrial Organization*, **9** (2), 303–313.
- COUTURE, V. (2016). Valuing the consumption benefits of urban density. Vancouver School of Economics, mimeograph.

- , DINGEL, J., GREEN, A., HANDBURY, J. and WILLIAMS, K. R. (2022). Measuring movement and social contact with smartphone data: a real-time application to covid-19. *Journal of Urban Economics*, **103**328.
- , DURANTON, G. and TURNER, M. A. (2018). Speed. *Review of Economics and Statistics*, **100** (4), 725–739.
- , GAUBERT, C., HANDBURY, J. and HURST, E. (2024). Income growth and the distributional effects of urban spatial sorting. *Review of Economic Studies*, **91** (2), 858–898.
- DAVIS, D. R., DINGEL, J. I., MONRAS, J. and MORALES, E. (2019). How segregated is urban consumption? *Journal of Political Economy*, **127** (4), 1684–1738.
- DEKLE, R., EATON, J. and KORTUM, S. (2007). Unbalanced trade. *American Economic Review*, **97** (2), 351–355.
- DELVENTHAL, M. and PARKHOMENKO, A. (2022). Spatial implications of telecommuting. *Available at SSRN 3746555*.
- DIAMOND, R. (2016). The determinants and welfare implications of us workers’ diverging location choices by skill: 1980-2000. *American Economic Review*, **106** (3), 479–524.
- DINGEL, J. and NEIMAN, B. (2020). How many jobs can be done at home? *Journal of Public Economics*, **189**, 104235.
- and TINTELNOT, F. (2023). Spatial economics for granular settings. *NBER working Paper*, **27287**.
- DOLFEN, P., EINAV, L., KLENOW, P. J., KLOPACK, B., LEVIN, J. D., LEVIN, L. and BEST, W. (2023). Assessing the gains from e-commerce. *American Economic Journal: Macroeconomics*, **15** (1), 342–70.
- DONALD, E., FUKUI, M. and MIYAUCHI, Y. (2024). Unpacking aggregate welfare in a spatial economy. *Working Paper*.
- EATON, B. C. and LIPSEY, R. G. (1982). An Economic Theory of Central Places. *The Economic Journal*, **92** (365), 56–72.
- EINAV, L., KLENOW, P. J., LEVIN, J. and MURCIANO-GOROFF, R. (2021). Customers and retail growth. *NBER Working Paper*, **29561**.
- FLORIDA, R. (2009). *Who’s Your City?: How the Creative Economy Is Making Where to Live the Most Important Decision of Your Life*. New York: Basic Books.
- FUJITA, M. and OGAWA, H. (1982). Multiple equilibria and structural transformation of non-monocentric urban configurations. *Regional Science and Urban Economics*, **12** (2), 161–196.
- GECHTER, M. and TSIVANIDIS, N. (2023). Spatial spillovers from urban redevelopments: Evidence from mumbai’s textile mills. University of California, Berkeley.

- GENTZKOW, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, **97** (3), 713–744.
- GLAESER, E. L., KOLKO, J. and SAIZ, A. (2001). Consumer city. *Journal of Economic Geography*, **1** (1), 27–50.
- GOKAN, T., KICHKO, S., MATHESON, J. and THISSE, J. (2022). *How the rise of teleworking will reshape labor markets and cities?* Tech. rep., Institute of Developing Economies, Japan External Trade Organization (JETRO).
- GORBACK, C. (2021). Ridesharing and the redistribution of economic activity. *Available at SSRN 4956559*.
- GUPTA, A., KONTOKOSTA, C. and VAN NIEUWERBURGH, S. (2022a). Take the Q Train: Value Capture of Public Infrastructure Projects. *Journal of Urban Economics*, **129**, 103422.
- , MITTAL, V. and VAN NIEUWERBURGH, S. (2022b). Work From Home and the Office Real Estate Apocalypse. *NBER Working Paper*, **30526**.
- HAUSMAN, N., SAMUELS, P., COHEN, M. and SASSON, R. (2023). Urban pull: The roles of amenities and employment. Harvard University, mimeograph.
- HOELZLEIN, M. (2023). Two-sided sorting and spatial inequality in cities. University of Notre Dame, mimeograph.
- HOTTMAN, C. J., REDDING, S. J. and WEINSTEIN, D. E. (2016). Quantifying the sources of firm heterogeneity. *Quarterly Journal of Economics*, **131** (3), 1291–1364.
- JIA, P. (2008). What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry. *Econometrica*, **76** (6), 1263–1316.
- KLOEK, T. and VAN DIJK, H. K. (1978). Bayesian Estimates of Equation System Parameters: An Application of Integration by Monte Carlo. *Econometrica*, **46** (1), 1–19.
- KOSTER, H. R. A., PASIDIS, I. and VAN OMMEREN, J. (2019). Shopping externalities and retail concentration: Evidence from dutch shopping streets. *Journal of Urban Economics*, **114** (August), 103194.
- KREINDLER, G. E. and MIYAUCHI, Y. (2023). Measuring commuting and economic activity inside cities with cell phone records. *The Review of Economics and Statistics*, **105** (4), 899–909.
- LEE, K. H. and TAN, B. J. (2024). Urban transit infrastructure and inequality. *The Review of Economics and Statistics*, pp. 1–46.
- LEONARDI, M. and MORETTI, E. (2023). The agglomeration of urban amenities: Evidence from milan restaurants. *American Economic Review: Insights*, **5** (2), 141–57.
- LUCAS, R. E. and ROSSI-HANSBERG, E. (2002). On the internal structure of cities. *Econometrica*, **70** (4), 1445–1476.

- MATSUYAMA, K. (1995). Complementarities and cumulative processes in models of monopolistic competition. *Journal of Economic Literature*, **33** (2), 701–729.
- McFADDEN, B. Y. D. (1989). A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica*, **57** (5), 995–1026.
- MELO, P. C., GRAHAM, D. J. and NOLAND, R. B. (2009). A meta-analysis of estimates of urban agglomeration economies. *Regional Science and Urban Economics*, **39** (3), 332–342.
- MILLS, E. S. (1967). An aggregative model of resource allocation in a metropolitan centre. *American Economic Review*, **57** (2), 197–210.
- MONTE, F., PORCHER, C. and ROSSI-HANSBERG, E. (2025). Remote work and city structure. *NBER Working Paper*, **31494**.
- , REDDING, S. and ROSSI-HANSBERG, E. (2018). Commuting, migration and local employment elasticities. *American Economic Review*, **108** (12), 3855–3890.
- MUTH, R. (1969). *Cities and Housing*. Chicago: University of Chicago Press.
- NEVO, A. and WONG, A. (2019). The elasticity of substitution between time and market goods: Evidence from the great recession. *International Economic Review*, **60** (1), 25–51.
- OBERFIELD, E., ROSSI-HANSBERG, E., SARTE, P.-D. and TRACHTER, N. (2024). Plants in space. *Journal of Political Economy*, **132** (3), 867–909.
- OH, R. and SEO, J. (2023). What causes agglomeration of services? theory and evidence from seoul. Yale University, mimeograph.
- OKUBO, T. (2022). Telework in the spread of covid-19. *Information Economics and Policy*, pp. 1–16.
- QIAN, F., ZHANG, Q. and ZHANG, X. (2024). Identifying agglomeration spillovers: Evidence from grocery store openings. Columbia University.
- REDDING, S. J. (2023). Quantitative urban models: From theory to data. *Journal of Economic Perspectives*, **37** (2), 75–98.
- (2025). Quantitative urban economics. In D. Donaldson and S. J. Redding (eds.), *Handbook of Regional and Urban Economics*, Amsterdam: Elsevier.
- and ROSSI-HANSBERG, E. (2017). Quantitative spatial models. *Annual Review of Economics*, **9**, 21–58.
- RELIHAN, L. E. (2022). Is Online Retail Killing Coffee Shops? Estimating the Winners and Losers of Online Retail using Customer Transaction Microdata. *Working Paper*.
- RUIZ, F. J. R., ATHEY, S. and BLEI, D. M. (2020). Shopper: A probabilistic model of consumer choice with substitutes or complements. *The Annals of Applied Statistics*, **14** (1), 1–27.
- SEVEREN, C. (2023). Commuting, labor, and housing market effects of mass transportation: Welfare and identification. *Review of Economics and Statistics*, **105** (5), 1073–1091.

- SHOAG, D. and VEUGER, S. (2018). Shops and the City: Evidence on Local Externalities and Local Government Policy from Big-Box Bankruptcies. *Review of Economics and Statistics*, **100** (3), 440–453.
- SMALL, K. and VERHOEF, E. (2007). *The Economics of Urban Transportation*. Routledge.
- STATISTICS BUREAU OF JAPAN (2016). 2016 economic census for business activity.
- STATISTICS BUREAU OF JAPAN (2019). Retail Price Survey.
- STATISTICS BUREAU OF JAPAN (2020). Family income and expenditure survey annual report (expenditure section), 2020.
- STATISTICS BUREAU OF JAPAN (2023). Report on internal migration based on the basic resident registration.
- SU, Y. (2022). Measuring the value of urban consumption amenities: A time-use approach. *Journal of Urban Economics*, **132**, 103495.
- THILL, J.-C. and THOMAS, I. (1987). Toward conceptualizing trip chaining: A review. *Geographical Analysis*, **19** (1), 1–17.
- THOMASSEN, Ø., SMITH, H., SEILER, S. and SCHIRALDI, P. (2017). Multi-Category Competition and Market Power: A Model of Supermarket Pricing. *American Economic Review*, **107** (8), 2308–2351.
- TRANSPORTATION RESEARCH BOARD (2006). *Commuting in America*. Washington DC: Department of Transportation.
- TSIVANIDIS, N. (2025). The aggregate and distributional effects of urban transit infrastructure: Evidence from bogotá’s transmilenio. *American Economic Review*, forthcoming.
- USHCHEV, P., SLOEV, I. and THISSE, J. F. (2015). Do We Go Shopping Downtown or in the ‘Burbs? *Journal of Urban Economics*, **85**, 1–15.
- VITALI, A. (2025). Consumer search and firm location: Theory and evidence from the garment sector in uganda. New York University, mimeo.
- ZÁRATE, R. D. (2021). Spatial misallocation, informality, and transit improvements: Evidence from mexico city. University of California, Berkeley.