

Rebuilding the Blocks: Firm Heterogeneity and Urban Retail

After Remote Work

Hoang-Anh Nguyen* Tuyetanh L. Tran†

October 10, 2026

Abstract

Urban neighborhoods differ not only in how many restaurants and shops they host, but also in the quality of those firms. We ask how spatial demand, competition, and spillovers shape the distribution of heterogeneous retail firms within and across locations, and how that distribution and consumer welfare respond to work-from-home adoption. Combining mobile-device GPS data with San Francisco administrative records, we document two patterns: firm-level visits vary more *within* Census block groups than across them, and consumers trip-chain, generating spillovers between neighbors. We then build and estimate a spatial equilibrium model in which consumers make nested location–firm choices and heterogeneous firms enter under incomplete information. Simulating work-from-home adoption, we find that city-center locations gain entrants despite losing foot traffic from local workers. Lower rents due to workers leaving, combined with these locations’ higher estimated inherent quality, outweigh the lost demand. In locations where entry falls, the welfare loss is smaller than counts alone suggest because the marginal firms are those with the lowest expected visits. Aggregate consumer surplus rises 2.7 to 8.0 percent across adoption levels, driven mainly by better access to retail.

JEL Codes: R10, R30, L11, D83

Keywords: agglomeration, firm entry, retail, urban amenities, GPS data, spatial spillovers, trip-chaining

*Tepper School of Business, Carnegie Mellon University; Email: hoangann@andrew.cmu.edu.

†Compass Lexecon; Email: ttran@compasslexecon.com.

1 Introduction

Urban locations exhibit striking variation in their retail business landscapes. Some neighborhoods host clusters of high-quality restaurants and shops that draw visitors from across the city. Others support only a few establishments serving primarily local residents. Understanding what drives these patterns has been central to urban economics for decades. A large literature studies why firms cluster in space and what determines the size of these clusters.¹ Yet this work has focused almost exclusively on firm counts, leaving the heterogeneity of firm quality within clusters largely unexplored.

In this paper, we study how spatial demand, competition, and spillovers shape the distribution of retail firms and their quality within and across locations, and how that distribution and consumer welfare respond to work-from-home adoption. Using mobile-device GPS data on visits to restaurants and shops in San Francisco, we document how firm quality varies within locations and estimate a spatial equilibrium model of consumer choice and firm entry. We then use the framework to solve for the new equilibrium of retail landscapes and consumption under remote work. Work-from-home is a major shock to urban retail: its widespread adoption after COVID-19 weakened demand in clusters that had served office workers while raising it in residential neighborhoods (Liu and Su, 2021; Ramani and Bloom, 2021; Qian and Su, 2025), prompting a wave of revitalization proposals.² Within each affected cluster, firms do not respond alike. A restaurant operating near break-even may exit when foot traffic moves against it. Better restaurants in the same cluster weather the shock and may even benefit from less competition. With enough exits, though, the cluster becomes less attractive to visitors, a negative spillover on the firms that remain. Our contribution is to bring this within-location heterogeneity into a quantitative spatial equilibrium framework

¹See, for instance, Ahlfeldt et al. (2015), Oh and Seo (2023), and Almagro and Domínguez-Iino (2025), among many others.

²See, for example, City and County of San Francisco (2023), San Francisco’s plan of nine strategies and nearly fifty initiatives for its downtown core.

with spillovers, where the quality of the marginal firms that respond to shocks matters for welfare.

Our analysis combines several data sources that allow us to observe both consumer mobility and the universe of non-tradable service firms in San Francisco in 2019.³ San Francisco experienced one of the largest shifts toward WFH in the United States and has since been the focus of policy efforts aimed at revitalizing the city.⁴ This makes it a natural setting to study how neighborhoods with different retail landscapes respond to urban shocks and regulations. Our core dataset comes from Veraset mobile-device GPS data, which records visits to points of interest (POIs) with timestamps, dwell times, and anonymized device identifiers. These data let us observe where consumers travel in the city, infer multi-stop trips, and measure trip-chaining behavior across locations and sectors, i.e., restaurants and shops. We then use sales tax data to show that GPS-based visits track business revenues, supporting their use as a proxy for consumer demand, from which we recover firm and location quality. Finally, we complement the GPS data with San Francisco administrative business records, which provide detailed information on firm entry, exit, expansion, and location changes. We estimate the model on pre-pandemic (2019) data and simulate WFH adoption and revitalization policies as counterfactuals.

We begin our analysis by documenting two empirical patterns, in firm heterogeneity and trip-chaining behavior, that motivate our structural model. First, firm-level performance, measured as GPS-based visits, varies substantially within locations. A variance decomposition of firm-level visits, run separately by sector, shows that around 75% of variation occurs within Census block groups (CBGs) rather than across them. Within-location dispersion is even more pronounced in popular locations. Treating firms within a location as symmetric, an assumption common in the quantitative spatial equilibrium literature, therefore

³We use *firm*, *establishment*, and *point of interest* interchangeably throughout; each refers to a single business location. [Section 2](#) sets out the full terminology.

⁴In 2021, roughly 35% of workers in the San Francisco–Oakland–Berkeley metropolitan area worked primarily from home, the highest share among U.S. metropolitan areas with at least one million residents and about twice the national rate of 17.9% ([U.S. Census Bureau, 2022](#)).

may abstract from a critical source of heterogeneity with important implications for local firm dynamics and welfare. Second, consumer trips often involve multiple stops, pointing to trip-chaining behavior that can generate demand spillovers across firms. Consumers make approximately 1.58 visits per trip on average, and 32% of consecutive stops occur within the same CBG. This suggests that high-quality firms can raise demand for their neighbors by drawing the first visit of a trip.

Building on these descriptive facts, we develop a spatial equilibrium model with endogenous firm entry and consumer choice. On the demand side, workers and residents make multi-stop trips, and where they go next depends on where they have just been. Each stop is a nested choice: consumers first choose a location (CBG) and then a specific firm within it, trading off travel costs against the variety and quality of firms. On the supply side, each location is endowed with a fixed number of potential entrants of heterogeneous quality. Potential entrants privately observe noisy signals about their own quality and simultaneously decide whether to enter, a one-shot entry game. The model features four key channels through which neighboring firms interact. First, trip-chaining and the nested demand structure generate demand spillovers across nearby firms. Second, firms compete directly for consumers. Third, scale economies capture productivity spillovers within locations, such as knowledge sharing and labor market pooling among same-sector firms. Finally, rents rise with the number of entrants and with the local resident and worker population, reflecting competition for limited commercial real estate.

Our model highlights how local firm heterogeneity shapes both the welfare effects of urban shocks and the entry response to them. All else equal, lower-quality firms respond first, because they operate closest to the entry margin. Since consumers choose these firms least often, shocks should have smaller welfare effects than a homogeneous-firm model would predict. Whether a shock leaves a location with more or fewer entrants is ambiguous, since the demand and rent channels move entry in opposite directions. Heterogeneity also matters for this balance: the quality of marginal firms varies across locations, so the spillovers they

carry differ accordingly, and with them the entry response.

We estimate the model in several steps. On the demand side, we use a nested logit framework in which consumers first choose a location to visit (a CBG) and then a specific firm within it. From observed location choices and trip-chaining behavior, we estimate travel costs, trip-chaining parameters, and location fixed effects. These fixed effects capture both location quality and the inclusive value of the firms located there. Conditional on location choice, observed firm choices then allow us to recover firm-level fixed effects, which identify firms’ own quality and location quality up to a scale. To fix the scale and separately identify location quality, we normalize the quality of branded firms to zero, assuming that they have the same quality wherever they locate. On the supply side, our goal is to estimate the parameters governing firm entry and local spillovers. We treat location quality, identified from the demand side, as endogenous: it reflects both exogenous location fundamentals and scale economies that rise with the number of firms in the location. The key identification insight is that fundamentals shape not only how many firms enter a location, but also the quality distribution within that location. This lets us disentangle exogenous fundamentals from the scale economies that entry generates. These fundamentals capture unobserved factors such as topography, microclimate, and historical infrastructure and land-use patterns, and we hold them fixed in counterfactuals.

Our estimates reveal substantial spatial frictions in consumer behavior. All else equal, visits are significantly more likely to occur near home, workplace, and previous trip destinations, confirming that proximity shapes both location choice and trip-chaining. We also find that more attractive locations exhibit greater dispersion in firm quality: better fundamentals typically support a wider range of firms rather than uniformly high-quality ones.

We then use the estimated model to simulate WFH adoption. We find that milder WFH adoption levels (25% and 50%) slightly decrease the total number of firms in the city, while more aggressive adoption (75%) increases it by about 1%. The effects vary across locations, and city-center areas see an increase in entrants despite office workers leaving. The main

channel is decreased rents in the worker-dense core, offsetting the loss of local demand there. City-center locations also have higher estimated inherent quality, independent of the firms located there, which lets them attract non-local visitors and sustain more firms despite losing foot traffic from local workers. Taken together, these results suggest that WFH adoption is less damaging to city-center retail than the decline in central business districts documented after COVID-19 might suggest (Ramani and Bloom, 2021; Rosenthal, Strange, and Urrego, 2022; Qian and Su, 2025). In locations where firm counts fall, heterogeneity means the firms that do not enter are those expecting the fewest visits, so consumer welfare falls by less than counts alone would suggest. Aggregate consumer welfare increases by 2.7 to 8.0 percent across WFH adoption levels, driven by increased retail consumption and lower travel costs. Our welfare measure abstracts from commuting savings, from any residential sorting the shock might induce, and from changes in landlords' rental income.

Related literature. This paper contributes to several strands within the broader urban and spatial economics literature.

First, we build on work studying how the spatial distribution of amenities shapes urban welfare and residential sorting. Ahlfeldt et al. (2015), Oh and Seo (2023), Almagro and Domínguez-Iino (2025), and Miyauchi, Nakajima, and Redding (2025) develop spatial equilibrium frameworks in which amenities are modeled through monopolistic competition within each location. These frameworks treat firms within a location as symmetric and focus on how the number and variety of amenities affect consumer welfare and residential choice. We contribute to this literature by allowing for heterogeneous firm quality within locations. This extension is motivated by our empirical finding that most of the variation in establishment performance occurs within rather than across locations. Our framework implies that firm quality, not just firm counts, matters for understanding agglomeration forces, policy impacts, and welfare outcomes.

Second, our paper relates to the literature on firm location choice and spatial competition.

Holmes (2011), Zhu and Singh (2009), and Seim (2006), among others, study store entry and location decisions. Existing work in this literature typically takes market characteristics and location quality as exogenous and focuses on location choices within a single sector or retail category. We contribute to this literature in several ways. First, we allow location quality to be endogenous through scale-economy spillovers generated by nearby firms. Second, we study equilibrium across multiple service sectors (e.g., restaurants and retail shops), allowing for spillovers both within and across amenity categories. Third, we endogenize demand linkages across locations through trip-chaining behavior, so that visits to one location affect demand in other locations. This richer framework allows us to study how demand shocks propagate through both firm entry and consumer mobility.

Third, we contribute to the emerging literature that uses GPS data to study urban development and mobility. Athey et al. (2018), Cao et al. (2024), Cook (2024), Couture, Dingel, et al. (2025), and Nguyen and Tran (2025) use smartphone location data to estimate heterogeneous preferences for amenities and measure travel patterns, with a primary focus on the demand side of spatial equilibrium models. Our contribution is to use GPS data to estimate spatial demand and then leverage these estimates, particularly the distribution of establishment-level fixed effects within locations, to identify supply-side objects, including productivity spillovers. This approach allows us to link observed mobility patterns to the underlying economic forces shaping firm entry and the spatial distribution of firm quality.

Finally, our paper relates to the growing literature examining how the shift to remote work affects cities. Recent work documents how the rise of WFH arrangements following COVID-19 has reshaped residential location choices, commercial real estate markets, and urban economic activity (Liu and Su, 2021; Ramani and Bloom, 2021; Rosenthal, Strange, and Urrego, 2022; Davis, Ghent, and Gregory, 2024; Qian and Su, 2025). We contribute to this literature by focusing on how service establishments respond to these demand shocks under competition and spatial spillovers. Our framework highlights that WFH trends can have heterogeneous effects across establishments of different quality: because entry adjusts at

the margin, the establishments that disappear where demand falls are those consumers visit least, while better establishments in the same location are insulated. This quality-dependent response has important policy implications, since firm counts alone understate how much of a neighborhood’s consumer value survives a shock.

The remainder of the paper is organized as follows. [Section 2](#) describes the data. [Section 3](#) presents the descriptive evidence motivating our model. [Section 4](#) develops the model. [Section 5](#) describes the estimation strategy. [Section 6](#) presents the estimation results. [Section 7](#) reports the work-from-home counterfactual. [Section 8](#) concludes.

2 Data

Our empirical analysis uses data from San Francisco. The city provides a particularly relevant setting for our research questions, which are motivated by the rise of remote work and subsequent rezoning efforts aimed at revitalizing metropolitan areas. Following the COVID-19 pandemic, San Francisco experienced one of the largest shifts toward work-from-home (WFH) in the United States ([U.S. Census Bureau, 2022](#)), substantially reducing foot traffic and demand for consumption amenities in many parts of the city. At the same time, the city has actively pursued rezoning and land-use reforms to revitalize its urban core. Together, these developments make San Francisco an ideal empirical setting for our research questions.

We construct our analysis data from several datasets. This section introduces our main data sources and presents their summary statistics. For estimation purposes, we use the data in June 2019.

Throughout the paper, we use the following terminology. A *sector* is a broad classification of non-tradable services (e.g., restaurants and shops); we use it interchangeably with *category*, and likewise use *non-tradable service* and *consumption amenity* interchangeably. A *location* is a small spatial unit in the city, which in our empirical analysis corresponds to a Census block group (CBG). A *destination* is a point of interest (POI), and we use *destination*,

establishment, *firm*, and *POI* interchangeably.⁵ Finally, *establishment quality* is the surplus an establishment generates for consumers when they visit and use its services; we use this interchangeably with *establishment performance* or simply *quality*.

2.1 Mobile GPS Data

Visit. Our mobile-device GPS data comes from Veraset, which provides coverage beginning in 2019 (Veraset, 2022).⁶ Veraset collects raw GPS pings and timestamps from anonymized devices and processes these sequences into visits to POIs. We use Veraset visit-level data, in which each visit records an anonymized device identifier, the device’s inferred home CBG, the visit timestamp and minimum dwell time, as well as destination characteristics such as geolocation, category (classified using NAICS codes), chain affiliation (brand), and on-site parking information.

We use sequences of visits made by the same device to proxy for trip-chaining behavior. We link consecutive visits that occur within a short time window, e.g., 3 hours, treating them as part of the same outing. This allows us to construct ordered pairs of visits, which we use to characterize trip-chaining patterns across sectors and locations in our analysis.

Consumer. The data allows us to construct a panel of consumers by linking visits made by the same device over time. For roughly 90% of devices, home locations are directly recorded in the data. By contrast, work locations are reported for less than 10% of devices, so we infer them from long visits during working hours (e.g., 6am to 8pm on weekdays). We follow Cook (2024) in constructing these measures, which yield work locations for roughly 50% of devices. This allows us to characterize each device by its residential and workplace locations and, in turn, construct the spatial distributions of residents and workers across the city. We validate these distributions using Census data and San Francisco travel survey data.

⁵A POI is a well-defined establishment, such as a restaurant or a flower shop.

⁶We access the Veraset data through Dewey (deweydata.io), an academic data platform (doi:10.82551/TP7E-VA02).

Establishment. We focus on visits to restaurants and retail establishments, which are the most common consumption amenities in San Francisco. Establishments are classified using NAICS codes, and we follow Cook (2024) in selecting the relevant codes.⁷

Summary statistics. Table 1 reports summary statistics for Veraset visit data at the CBG and establishment levels. Our sample covers roughly 500 CBGs in San Francisco, containing a total of about 4,000 restaurants and 3,500 retail shops. The median CBG has 4 restaurants and 3 shops. There is substantial heterogeneity across CBGs, with establishment counts ranging from 1 to over 300. Restaurant visits account for roughly 60 percent of total visits to POIs in the two sectors of interest (restaurants and retail shops) during June 2019. Chain-affiliated establishments are relatively rare, with a CBG containing about one branded restaurant or shop on average. Restaurants also exhibit longer dwell times, consistent with the nature of dining versus shopping trips.

⁷See Table A.1 in Cook (2024) for more details.

Table 1: Summary Statistics – Mobile GPS Data

	Mean	P25	Median	P75
<i>Panel A: CBG-level</i>				
Number of establishments	15.29	2.50	8.00	19.00
Restaurants	8.48	1.00	4.00	11.00
Branded	0.78	0.00	0.00	1.00
Shops	6.81	1.00	3.00	7.50
Branded	1.07	0.00	0.00	1.00
Total visits	3,367	492	1,271	3,432
Unique visitors	1,755	256	695	1,890
<i>Panel B: Establishment-level (Restaurants)</i>				
Visits	262.7	81.0	164.0	289.0
Unique visitors	183.5	58.0	114.0	200.0
Avg. dwell time (min.)	136.2	73.3	108.9	163.9
<i>Panel C: Establishment-level (Shops)</i>				
Visits	167.2	17.0	68.0	178.2
Unique visitors	112.0	14.0	47.0	111.0
Avg. dwell time (min.)	92.7	33.8	68.3	119.7

Notes: This table reports summary statistics for Veraset visit data at the CBG level (Panel A) and establishment level for restaurants (Panel B) and shops (Panel C). Visits and unique visitors are measured over our analysis period of June 2019 from a sample of mobile devices (Veraset reports coverage of roughly 6% of the U.S. population). Dwell time is measured in minutes.

2.2 Business Sales Tax Data

We use business sales tax data from the San Francisco Controller’s Office. The dataset records sales tax at the tax block group level, groups of Census blocks defined for tax reporting purposes, and is available at the quarter level for calendar years beginning in 2015. We use these data to validate our mobile GPS measure of visits by examining the relationship between visit counts and realized business revenues. In our model, a visit corresponds to consumption of the service, whereas in the data a GPS-based visit need not imply a purchase. A strong positive relationship between visits and revenues therefore supports the use of GPS-based visits as a proxy for realized consumer demand and, in turn, for establishment quality and attractiveness.

Summary statistics. [Table 2](#) reports summary statistics for sales tax data at the tax block group level. Our sample contains about 600 tax block groups. The median tax block group has 5 establishments, generating roughly \$19,000 in sales tax in the second quarter of 2019. There is substantial variation in sales tax across tax block groups: the interquartile range spans approximately \$7,000 to \$45,000.

Table 2: Summary Statistics – Sales Tax Data

	Mean	P25	Median	P75
Sales tax (\$)	45,410	6,669	19,470	43,620
Total visits	1,072	424	756	1,294
Number of establishments	6.71	3.00	5.00	8.00
Restaurants	2.79	1.00	2.00	4.00
Shops	2.89	1.00	2.00	3.00

Notes: This table reports summary statistics for sales tax data at the tax block group level for the second quarter of 2019. Sales tax is measured in dollars. Total visits and establishment counts are constructed from matching tax block groups to Veraset visit data. Visits are measured over our analysis period of June 2019 from a sample of mobile devices (Veraset reports coverage of roughly 6% of the U.S. population). Restaurants and shops are classified using NAICS codes.

2.3 Business Registry Data

We also supplement the GPS data with business registration records from the San Francisco Treasurer and Tax Collector’s Office. These records provide information on all registered businesses in the city, including establishment names, owners, addresses, business types, registration dates, business opening and closing dates, and the locations of different branches, along with their opening and closing histories. The dataset includes all businesses active as of 2014 and all businesses opened thereafter, including those that later became inactive. We use this dataset to verify establishment characteristics, fill in missing NAICS codes, and ensure that our sample captures the universe of service businesses operating during our analysis period.

3 Descriptive Evidence

We document descriptive statistics based on the three datasets described in [Section 2](#) and consider the two non-tradable service sectors central to our analysis, namely restaurants and retail shops. For our descriptive analysis, we further distinguish full-service restaurants from other restaurants (e.g., coffee shops and fast-food restaurants) within the restaurant sector.

Validation of GPS-Based Visits. We validate the mobile GPS measure of visits using business sales tax data at the tax block group level. Specifically, we examine whether tax block groups receiving more GPS-based visits also generate higher realized business revenues, as proxied by sales tax. A positive relationship between visits and revenues supports the use of GPS-based visits as a proxy for realized consumer demand and, in turn, for establishment quality and attractiveness.

We estimate the following regression at the tax block group level

$$\log(\text{SalesTax}_g) = \beta_1 \log(\text{Visits}_g) + \beta_2 \text{Estab}_g + \alpha_{n(g)} + \varepsilon_g \quad (1)$$

where g indexes tax block groups, SalesTax_g denotes the total sales tax, Visits_g is the number of GPS-based visits, and Estab_g is the number of establishments in the tax block group. We include neighborhood fixed effects, $\alpha_{n(g)}$, for the neighborhood in which tax block group g is located.⁸

[Table 3](#) reports the result. We find a positive relationship between GPS-based visits and realized business revenues: a 1% increase in visits is associated with a 0.26% increase in sales tax, controlling for the number and composition of establishments as well as neighborhood fixed effects. The coefficient is statistically significant at the 1% level. This result supports the use of GPS-based visits as a proxy for establishment quality and attractiveness in the

⁸In this paper, we define neighborhood based on the Analysis Neighborhood definition provided by the San Francisco Mayor’s Office of Housing and Community Development.

analysis that follows.

Table 3: GPS-Based Visits and Sales Tax

Log sales tax	Estimate	Std. Error
Log visits	0.261***	0.078
Share full-service restaurants	0.000	0.001
Share other restaurants	0.001	0.003
Share shops	0.000	0.001
Number of establishments	0.044**	0.014
Neighborhood FE	Yes	

Notes: This table reports estimates from a regression of log sales tax on log GPS-based visits at the tax block group level. The regression controls for the number and composition of establishments (shares of full-service restaurants, other restaurants, and shops). The dependent variable is the log of total sales tax in the second quarter of 2019, which covers our analysis period of June 2019. The regression includes neighborhood fixed effects. Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. The regression also uses Veraset visit data for June 2019.

Within-Location Firm Quality Variation. We document substantial heterogeneity in establishment performance within locations. Using GPS-based visits as a proxy for establishment quality, we decompose the variation of establishment quality into within-location and between-location components.

Formally, for each sector $s \in \{\text{full-service restaurant, other restaurant, shop}\}$, let X_k denote the number of visits to establishment k , and let $j(k)$ denote the CBG in which

establishment k is located. We decompose the variance of visits as follows⁹

$$\text{Var}(X_k) = \mathbb{E} [\text{Var}(X_k \mid j(k))] + \text{Var} (\mathbb{E}[X_k \mid j(k)]) \quad (2)$$

where the first term captures within-location variation and the second term captures between-location variation.

[Figure 1](#) presents this decomposition by sector. Across all categories, the majority of the variation in visits occurs within locations rather than between locations: the within-CBG share is roughly 78% for full-service restaurants, 72% for other restaurants, and 76% for shops. This pattern indicates that establishments within the same sector and CBG can attract substantially different levels of consumer demand.

⁹Details of the derivation are provided in [Section B.1](#).




Figure 1: Decomposition of Visit Variation by Category

Notes: This figure decomposes the total variance of log visits across establishments into within-CBG and between-CBG components, separately for full-service restaurants, other restaurants, and retail shops. The within-CBG component captures variation in visits across establishments located in the same CBG, while the between-CBG component captures variation in average visits across CBGs. The figure is constructed using Veraset data for June 2019.

Figure 2 reports the variance decomposition for full-service restaurants across CBGs, grouped by overall visit levels. A clear pattern emerges across quantiles: as we move from less to more popular locations, the within-CBG share of total variation increases, while the between-CBG share declines. In other words, in more popular locations, which likely reflect better local fundamentals, variation in visits is increasingly driven by differences across establishments within the same CBG rather than across CBGs. Even in less popular locations, within-location heterogeneity remains significant.



Figure 2: Decomposition of Visit Variance by CBG Visit Quantile (Full-Service Restaurants)

Notes: This figure reports the variance decomposition of establishment visits to full-service restaurants into within-CBG and between-CBG components, separately by CBG visit quantile. CBGs are grouped by their overall visit levels, so higher quantiles correspond to more popular locations. The figure is constructed using Veraset data for June 2019.

We further examine how the relative importance of within-location and between-location variation differs across neighborhoods.¹⁰ For each neighborhood, we compute the ratio of within-CBG to between-CBG variation in establishment visits and visualize the result in [Figure 3](#). The figure shows that in most neighborhoods, within-location variation accounts for a substantial share of the total variation in establishment visits and exceeds between-location variation. Neighborhoods where between-location variation is relatively more important tend to be less popular, more residential areas, where the majority of CBGs contain no more than one establishment, leaving less scope for within-location variation.

¹⁰A location, such as a CBG, is a smaller spatial unit that lies within a neighborhood.




Figure 3: Between-to-Within Visit Variance Ratio by Neighborhood (Full-Service Restaurants)

Notes: This figure shows the ratio of within-CBG to between-CBG variation in visits to full-service restaurants for each neighborhood. A ratio greater than one indicates that there is relatively more between-location variation than within-location variation. Neighborhoods are shaded according to the magnitude of this ratio. Neighborhoods with substantial missing values are omitted because there is insufficient establishment visit data to compute the ratio. The figure is constructed using Veraset visit data for June 2019.

Together, these patterns suggest that differences in establishment performance cannot be explained solely by location-level characteristics. Instead, establishments within the same location exhibit substantial heterogeneity in the demand they attract, which we use as a proxy for their quality. This motivates our structural model, which allows for heterogeneous firm quality within locations rather than treating establishments in the same location as symmetric. Moreover, the increase in within-location variation across location quantiles is central to our identification strategy: locations with better underlying fundamentals also ex-

hibit greater within-location dispersion in establishment performance. This variation allows us to separately identify location fundamentals.

Evidence of Trip-Chaining. We provide evidence of trip-chaining behavior in consumers’ mobility patterns. On average, consumers make approximately 1.58 visits per trip in our data. This estimate is consistent with existing evidence in the literature. For example, Oh and Seo (2023) reports an average of 1.72 visits per travel.

Table 4 reports transition probabilities across categories within a trip. The result shows that trip-chaining occurs both within and across different amenity types. Trip-chaining also occurs across locations. Table 5 reports whether consecutive visits occur within the same location or across different locations. We find that 32% of linked visits occur within the same CBG, while 68% occur across different CBGs.

Table 4: Trip-Chaining by Category

	Full-service restaurants	Other restaurants	Shops
Full-service restaurants	0.507	0.204	0.289
Other restaurants	0.402	0.300	0.298
Shops	0.377	0.189	0.435

Notes: This table reports transition probabilities between amenity categories within the same trip. Each row indicates the category of the first visit and each column indicates the category of the subsequent visit. Entries represent the probability that a visit to one category is followed by a visit to another category within the same trip. The probabilities in each row sum to one. Trips are constructed by linking consecutive visits by the same device within a three-hour window. The table is constructed using Veraset visit data for June 2019.

Table 5: Trip-Chaining by Location

Transition type	Share (%)
Across CBG	68
Within CBG	32

Notes: This table reports the share of consecutive visit pairs that occur within the same CBG versus across different CBGs. Trips are constructed by linking consecutive visits by the same device within a three-hour window. The table is constructed using Veraset visit data for June 2019.

Taken together, these patterns provide clear evidence of trip-chaining behavior across both sectors and locations. This behavior can generate demand spillovers across establishments, as visits to one amenity may increase demand for nearby amenities or other consumption categories within the same trip. Motivated by this evidence, our structural model allows consumers to trip-chain.

Local Visit Shares and Nested Location Choice Structure. [Section 4](#) proposes a nested logit demand model in which consumers first choose a location j and then choose an establishment k within that location. This nested structure implies a simple and testable restriction on the share of local visitors to an establishment. We can write the share of local visitors to establishment k in location j as

$$\mathbb{P}_{\text{local}|k} = \frac{\mathbb{P}_{j|\text{local}} N_{\text{local}}}{\mathbb{P}_{j|\text{local}} N_{\text{local}} + \mathbb{P}_{j|\text{nonlocal}} N_{\text{nonlocal}}} \quad (3)$$

The nesting structure implies that the share of local visitors depends only on location characteristics and not on establishment-level attributes. We test this implication by examining whether establishments that attract more visits also have systematically different shares of local visitors. Specifically, we regress the share of visitors whose home or work location is near an establishment on the log of visits to that establishment. [Table 6](#) reports the results. When we control for location fixed effects, the coefficient on log visits is statistically insignificant. In contrast, when we do not control for CBG fixed effects, establishments with more visits tend to have lower shares of local residents or higher shares of nearby workers.¹¹ Overall, these results are consistent with the nesting implication: once location is held fixed, establishment-level demand does not affect the share of local visitors. This evidence suggests that the nesting structure is a reasonable model choice.

¹¹The opposite signs are consistent with higher-visit destinations typically being located in commercial areas with more office buildings. These areas have fewer local residents but more nearby workers.

Table 6: Local Visitor Share and Establishment Visit

Local visitor share	Resident		Worker	
	(1)	(2)	(3)	(4)
Log visits	-0.213 (0.304)	-1.262*** (0.290)	0.257 (0.275)	1.554*** (0.284)
Category (Sector) FE	Yes	Yes	Yes	Yes
Location FE	Yes	No	Yes	No

Notes: This table reports regressions of the share of local visitors on establishment visit counts. The dependent variable is the share of visitors whose home location (Columns (1) and (2)) or work location (Columns (3) and (4)) is located in the same CBG with the establishment. The main regressor is the log of GPS-based visits to the establishment. Columns (1) and (3) include location fixed effects, while columns (2) and (4) omit them. All specifications include category fixed effects. Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. The table is constructed using Veraset visit data for June 2019.

4 Model

4.1 Setup

We partition the city into J locations and characterize individuals by their home and work locations. Let $j \in \mathcal{J} \equiv \{1, 2, \dots, J\}$ index locations, which differ in underlying fundamentals, such as topography, proximity to open space and nature, and relative distances to one another, as well as in the number and composition of residents, workers, and establishments.

For example, in our empirical context, a location is a CBG. There are N_{hw} people living in location h and working in location w . The city has S non-tradable service sectors, indexed by $s \in \mathcal{S} \equiv \{1, 2, \dots, S\}$, and N_{sj} establishments in sector-location sj .

Demand for consumption amenities comes from residents and workers across the city, and the distribution of consumer origins (i.e., their home and work locations) is taken as given. Each consumer chooses a destination to visit, i.e., an establishment $k \in \mathcal{K}_{sj}$ where $\mathcal{K}_{sj}$ denotes the set of establishments in sector s located in location j . Consumers may trip-chain, making a sequence of visits to multiple destinations in a single trip. On the supply side, heterogeneous potential entrants (i.e., establishments) choose whether to enter a given location, conditional on location fundamentals as well as their expectations about demand, competition, spillovers, and rents.

4.2 Demand

Consumer i , with home location $h(i)$ and work location $w(i)$, exogenously draws a non-tradable service sector s to visit and then chooses a destination. The consumer first selects a location and then chooses an establishment within that location. This two-stage decision is represented by a two-level nested logit (Cardell, 1997; McFadden, 1978), where each sector-location defines a nest, and establishments within that sector-location are the alternatives within the nest. Upon visiting the first destination, the consumer may trip-chain and visit additional destinations, potentially in sectors other than the one initially drawn.

Stage 2: Establishment Choice Conditional on Location. Conditional on drawing sector s and selecting location j , consumer i then chooses an establishment $k \in \mathcal{K}_{sj}$. Because the data contains few establishment-level observables beyond name, NAICS code, and geolocation, we model the mean utility of an establishment (conditional on s and j) using establishment fixed effects rather than as a function of observed characteristics as in standard BLP-style models. Let δ_{sjk} denote the establishment fixed effect (i.e., mean utility) for

establishment k in sector s and location j . We also interpret δ_{sjk} as establishment quality.

The value of establishments in sector-location sj is the expected utility from the within-location establishment choice, conditional on drawing sector s and visiting location j . Under the nested logit structure, this value has a closed-form expression:

$$\text{Estab}_{sj} = \rho \log \left[\sum_{k=1}^{N_{sj}} \exp(\delta_{sjk}/\rho) \right] \quad (4)$$

where $\rho > 0$ is a nesting parameter that links establishment and location choices. A higher ρ implies weaker substitution across establishments within the same location. The market share of establishment k conditional on visiting sector-location sj is

$$s_{k|sj} = \frac{\exp(\delta_{sjk}/\rho)}{\sum_{k'=1}^{N_{sj}} \exp(\delta_{sjk'}/\rho)} \quad (5)$$

Stage 1: Location Choice Conditional on Sector Draw. Consumer i draws a non-tradable service sector s and chooses a location $j \in \mathcal{J}$ to visit. The mean utility from visiting location j for a consumer who draws sector s and has home location h and work location w is

$$\delta_{sjhw} = \gamma_{shw} + \beta_{\text{home}} \text{Dist}_{hj} + \beta_{\text{work}} \text{Dist}_{wj} + \underbrace{\beta_{TC} \text{TripChain}_{sj}}_{\gamma_{sj}} + \text{Estab}_{sj} + \xi_{sj} \quad (6)$$

where γ_{shw} represents sector-home-work fixed effects, and Dist_{hj} and Dist_{wj} denote the distances from j to the consumer's home and work locations, respectively. γ_{sj} represents sector-location fixed effects that can be decomposed into three components: the expected value of trip-chaining, TripChain_{sj} (discussed below); the inclusive value of establishments in sector-location sj , Estab_{sj} ; and unobserved preferences for location j conditional on sector s , ξ_{sj} . We also allow for an outside option, denoted $j = 0$, which corresponds to not making a trip within the city. Choosing this option yields normalized utility zero.

The location market share (probability of visiting sector-location sj given origin hw) is

$$s_{j|shw} = \frac{\exp(\delta_{sjhw})}{1 + \sum_{j' \in \mathcal{J}} \exp(\delta_{sj'hw})} \quad (7)$$

Trip-Chaining. The value of trip-chaining for the consumer is

$$\text{TripChain}_{sj} = \sum_{s' \in \mathcal{S}} f(s'|s) \log \left[1 + \sum_{j' \in \mathcal{J}} \exp \left(\frac{\beta_{\text{cont}} \text{Dist}_{jj'} + \tilde{\gamma}_{js'} + \text{Estab}_{s'j'} + \xi_{s'j'}}{\beta_{TC}} \right) \right] \quad (8)$$

where $f(s'|s)$ denotes the exogenous probability that the second visit is to sector s' , conditional on the first visit being to sector s , and $\tilde{\gamma}_{js'}$ represents the fixed effects of drawing sector s' after visiting location j . For simplicity, we assume that each trip consists of at most two visits. If the consumer does not make a second visit within the same trip (i.e., no trip-chaining is realized), we normalize the utility of not making a second visit to zero, similar to the outside option of not visiting the city.

4.3 Supply

In this section, we model establishments' strategic entry decisions. This allows us to characterize how the number and composition of firms in each sector-location, as well as sector-location unobservables, are determined endogenously. For notational simplicity, we suppress the sector index s throughout this section and present the model for a single sector. The multi-sector extension is straightforward: sectors interact through consumer trip-chaining, which is estimated on the demand side but abstracted from here.

4.3.1 Location Quality and Spillovers

Each sector-location j is characterized by an *inherent* quality $\tilde{\delta}_j^{\text{loc}}$, which captures exogenous location fundamentals such as topography, microclimate, historical infrastructure, and land-use patterns. Location quality as perceived by consumers and entrants is augmented by

agglomeration spillovers from entering firms:

$$\xi_j = \tilde{\delta}_j^{\text{loc}} + \alpha N_j \quad (9)$$

where $N_j = \sum_{k=1}^{N_j^{\text{pot}}} d_{jk}$ is the equilibrium number of entrants at location j , $d_{jk} \in \{0, 1\}$ is the entry decision of firm k , and $\alpha \geq 0$ is the agglomeration parameter. The term αN_j captures economies of scale, i.e., productivity spillovers among co-located firms arising from, for example, knowledge sharing, labor market pooling, and consumer traffic externalities.

4.3.2 Profit Function

The profit of firm k entering location j depends on the consumer visits it attracts, the local commercial rent, and a fixed operating cost. We specify a log-linear profit function:

$$\pi_{jk} = \beta \log(\text{visits}_{jk}) - \gamma_0 - \gamma_1 \log(w_j) - \gamma_2 \log(r_j) - \gamma_3 \log(N_j) \quad (10)$$

where:

- visits_{jk} denotes the number of consumer visits to firm k at location j , determined by the demand model given the set of available locations and firms;
- γ_0 is a fixed cost;¹²
- $\gamma_1 \log(w_j)$ captures rent driven by local worker density;
- $\gamma_2 \log(r_j)$ captures rent driven by local resident density;
- $\gamma_3 \log(N_j)$ captures rent driven by local business density.¹³

¹²Multiplying the home-work matrix by any constant $c > 0$ shifts $\log(\text{visits})$, $\log(w_j)$, and $\log(r_j)$ each by $\log c$, leaving β , γ_1 , and γ_2 unchanged but shifting γ_0 by $(\beta - \gamma_1 - \gamma_2) \log c$. The fixed cost γ_0 therefore absorbs both the true fixed operating cost and any level adjustments arising from sample coverage and home-work matrix scaling, and should not be interpreted as a dollar-denominated cost.

¹³Couture, Gaubert, et al. (2024) assume rent has function N_j^γ , where γ is the inverse supply elasticity of real estate.

The parameter γ_3 governs the rent competition from local entrants and is fixed for normalization in our estimation procedure. The log-linear specification implies diminishing returns to consumer traffic: each additional visitor raises profit by β/visits , a decreasing function of the current visit level. Consumer visits are determined by the nested logit demand model:

$$\text{visits}_{jk} = \sum_{h,w} N_{hw} \cdot s_{j|shw} \cdot s_{k|sj} \quad (11)$$

where $s_{j|shw}$ is the location choice share (Equation (7)) and $s_{k|sj}$ is the establishment choice share (Equation (5)).

4.3.3 Firm Quality and Private Signals

Each potential entrant k at location j draws a true quality

$$\delta_{jk} \sim \mathcal{N}(\mu, \sigma_\delta^2) \quad (12)$$

independently across firms and locations, where μ is the mean quality of the potential entrant pool and σ_δ governs dispersion around that mean. Before the entry decision, the firm observes only a noisy private signal of its own quality:

$$\tilde{\delta}_{jk} = \delta_{jk} + \nu_{jk}, \quad \nu_{jk} \sim \mathcal{N}(0, \sigma_\nu^2) \quad (13)$$

where ν_{jk} is an idiosyncratic entry shock, independent of δ_{jk} and independent across firms. The parameter σ_δ governs the dispersion of true firm quality across potential entrants, while σ_ν governs the noisiness of each firm's private signal. The signal has unconditional variance $\tilde{\sigma}^2 \equiv \sigma_\delta^2 + \sigma_\nu^2$. Upon entry, the firm's true quality δ_{jk} is revealed and determines its consumer visits.

All potential entrants commonly observe before entry decisions are made:

- the home-work distribution $\{N_{hw}\}$;

- the vector of inherent location qualities $\{\tilde{\delta}_j^{\text{loc}}\}_{j \in \mathcal{J}}$;
- the vector of potential entrant masses $\{N_j^{\text{pot}}\}_{j \in \mathcal{J}}$;
- the profit function $\pi(\cdot)$ and its parameters;
- the distributions $F_\delta = \mathcal{N}(\mu, \sigma_\delta^2)$ and $F_\nu = \mathcal{N}(0, \sigma_\nu^2)$.

4.3.4 Entry Equilibrium

Each potential entrant k decides whether to enter based on its private signal $\tilde{\delta}_{jk}$, taking as given its beliefs about the entry decisions of all other potential entrants. Note that expected profit is strictly increasing in the private signal: higher signal implies higher expected quality, higher expected visits, and higher expected profit. Therefore, entering if and only if the signal exceeds a threshold is the best response to any strategy profile of others. We focus on such threshold equilibria.

Definition 1 (Equilibrium). *A vector of thresholds $\tau^* = \{\tau_j^*\}_{j \in \mathcal{J}}$ constitutes a symmetric Bayesian-Nash equilibrium if, for every location j , the marginal entrant at threshold τ_j^* earns zero expected profit:*

$$\mathbb{E}\left[\pi_{jk} \mid \tilde{\delta}_{jk} = \tau_j^*, \{d_{j'k'} = \mathbf{1}[\tilde{\delta}_{j'k'} \geq \tau_{j'}^*]\}_{j' \in \mathcal{J}}\right] = 0 \quad \forall j \in \mathcal{J} \quad (14)$$

where the expectation is taken over the draws of all other firms' qualities and entry decisions, which are determined by the threshold strategies $\{\tau_{j'}^*\}_{j' \neq j}$.

The entry rule is:

$$d_{jk} = \begin{cases} 1 & \text{if } \tilde{\delta}_{jk} \geq \tau_j^* \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

Proposition 1 (Existence of Equilibrium). *There exists a symmetric Bayesian-Nash equilibrium in threshold strategies.*

Proof. Define the operator $T : [0, 1]^J \rightarrow [0, 1]^J$ acting on the vector of entry probabilities $p_j = 1 - \Phi(\tau_j^*/\tilde{\sigma})$ by:

$$T_j(p) = \max(0, \min(1, p_j + \lambda \cdot \mathbb{E}[\pi_j \mid p])) \quad (16)$$

where $\mathbb{E}[\pi_j \mid p]$ is the expected profit of the marginal firm at location j given that all locations play entry probabilities p , and $\lambda > 0$ is a step size. T is a continuous mapping from the compact convex set $[0, 1]^J$ to itself. By Brouwer's fixed-point theorem, T has a fixed point $p^* \in [0, 1]^J$. At this fixed point, $\mathbb{E}[\pi_j \mid p^*] = 0$ for all j with $p_j^* \in (0, 1)$, which corresponds to interior equilibrium thresholds $\tau_j^* = \tilde{\sigma} \cdot \Phi^{-1}(1 - p_j^*)$. $\square$

Remark 1. *The equilibrium need not be unique: multiple threshold vectors may satisfy the zero-profit conditions when agglomeration economies ($\alpha > 0$) create strategic complementarities across locations. In our estimation, we focus on the equilibrium selected by the iterative threshold-updating algorithm initialized from the observed entry rates, which we interpret as the empirically relevant equilibrium.*

4.4 Market Equilibrium

Given the entry equilibrium and realized establishment fixed effects, the post-entry market equilibrium consists of the realized:

1. Entry outcomes $\mathbf{d} = \{d_{sjk}\}_{s \in \mathcal{S}, j \in \mathcal{J}, k \in \mathcal{K}_{sj}}$, given Equation (15) as well as the draw of firm qualities and private signals;
2. Firm profits $\boldsymbol{\pi} = \{\pi_{sjk}\}_{s \in \mathcal{S}, j \in \mathcal{J}, k \in \mathcal{K}_{sj}}$, given by Equation (10);
3. Sector-location unobservables $\boldsymbol{\xi}$, given by Equation (9);
4. Market shares and trip-chaining shares, given by Equations (5), (7) and (8).

The sector-location unobservable ξ_{sj} , which plays a central role in the model, reflects both exogenous location fundamentals $\tilde{\delta}_{sj}^{\text{loc}}$ and endogenous agglomeration spillovers. The

payoff of a firm at sector-location sj depends on its own quality δ_{sjk} , the qualities of other firms δ_{-k} (whether competing or complementary), the sector-location unobservables ξ , and local commercial rent c_j :

$$\pi_{sjk} = \pi(\delta_{sjk}, \delta_{-k}, \xi, c_j; \alpha_\pi) \quad (17)$$

where α_π denotes the vector of payoff parameters. Since we do not observe commercial rents at the location level, we follow Couture, Gaubert, et al. (2024) and Almagro and Domínguez-Iino (2025) and parameterize c_j as a function of observables that proxy for local commercial demand and land-use intensity. In particular, we parameterize c_j as a function of the distribution of consumer origins and establishment densities by sector-location (Equation (10)).

5 Identification and Estimation

This section discusses how to identify and estimate the model in Section 4. We first estimate demand for consumption amenities to recover travel cost parameters, sector-location unobservables, and establishment fixed effects. We then estimate the supply side, exploiting the systematic variation in within-location distributions of establishment qualities to separately identify sector-location fundamentals and supply-side spillovers.

5.1 Demand

We estimate the demand side in two main steps. First, we estimate travel cost parameters β 's, including trip-chaining, and the decomposition of sector-location fixed effects γ_{sj} . Second, we separately identify sector-location unobservables ξ_{sj} and (scaled) firm fixed effects δ_{sjk} .

5.1.1 Travel Costs, Trip-Chaining, and Sector-Location Fixed Effects

First, we derive the regression equations corresponding to the first and second visits within a trip:¹⁴

$$\log N_{sjhw} = \gamma_{shw} + \beta_{\text{home}} \text{Dist}_{hj} + \beta_{\text{work}} \text{Dist}_{wj} + \gamma_{sj} + \varepsilon_{sjhw} \quad (18)$$

$$\log N_{js'j'} = \frac{\beta_{\text{cont}}}{\beta_{TC}} \text{Dist}_{jj'} + \hat{\gamma}_{js'} + \tilde{\gamma}_{s'j'} + \varepsilon_{js'j'} \quad (19)$$

Here N_{sjhw} is the number of consumers visiting sector-location sj with home-work origin hw and $N_{js'j'}$ is the number of consumers who visit sector-location $s'j'$ after initially visiting j during the trip. The ε terms denote regression errors. We also have

$$\tilde{\gamma}_{s'j'} = \frac{1}{\beta_{TC}} [\text{Estab}_{s'j'} + \xi_{s'j'}] \quad (20)$$

Estimating these regressions yields β_{home} , β_{work} , $\frac{\beta_{\text{cont}}}{\beta_{TC}}$, and the fixed effects γ 's, $\hat{\gamma}$'s, and $\tilde{\gamma}$'s.

Next, let $\bar{\gamma}_{sj}$ denote the aggregate sector-location effects, i.e., the sum of the sector-location unobservable ξ_{sj} and the establishment inclusive value Estab_{sj} :

$$\bar{\gamma}_{sj} = \text{Estab}_{sj} + \xi_{sj} = \gamma_{sj} - \beta_{TC} \text{TripChain}_{sj} \quad (21)$$

We can rewrite γ_{sj} and $\tilde{\gamma}_{sj}$ as functions of $\bar{\gamma}_{sj}$:

$$\gamma_{sj} = \bar{\gamma}_{sj} + \beta_{TC} \sum_{s' \in \mathcal{S}} f(s'|s) \log \left[1 + \sum_{j' \in \mathcal{J}} \exp \left(\frac{\beta_{\text{cont}} \text{Dist}_{jj'} + \hat{\gamma}_{js'} + \tilde{\gamma}_{s'j'}}{\beta_{TC}} \right) \right] \quad (22)$$

$$\tilde{\gamma}_{sj} = \frac{\bar{\gamma}_{sj}}{\beta_{TC}} \quad (23)$$

These equations characterize a system in unknowns $\{\bar{\gamma}_{sj}\}_{s \in \mathcal{S}, j \in \mathcal{J}}$ and β_{TC} . Given our es-

¹⁴We abuse notation here and write $\tilde{\gamma}_{js'}$ to denote the scaled object $\frac{\gamma_{js'}}{\beta_{TC}}$; the unscaled $\tilde{\gamma}_{js'}$ is defined in [Section 4.2](#). Details of the derivations are provided in [Sections B.2](#) and [B.3](#).

timates of $\{\gamma_{sj}, \hat{\gamma}_{js'}, \tilde{\gamma}_{sj}\}_{s \in \mathcal{S}, j \in \mathcal{J}}$ and $\frac{\beta_{\text{cont}}}{\beta_{TC}}$, along with the empirical transition probabilities $f(s'|s)$ and pairwise location distances, we recover $\{\bar{\gamma}_{sj}\}_{s \in \mathcal{S}, j \in \mathcal{J}}$ and β_{TC} as the values that best fit the system above.

Intuitively, the variation in visits across locations identifies aggregate sector-location effects, i.e., the sum of sector-location unobservables and the establishment inclusive value within a sector-location. We also exploit the variation in observed trip-chaining behavior to estimate the trip-chaining effects.

5.1.2 Sector-Location Unobservables and Establishment Fixed Effects

In this section, we exploit the variation in establishment choices, conditional on drawn sectors and chosen locations, to identify establishment fixed effects (scaled by the nesting parameter ρ).

Assume that, in each sector-location, there exists a reference establishment 0 where $\delta_{sj0} = \delta_{s0}$ with δ_{s0} being known. For example, we can use a chain establishment (e.g., McDonald's) as the reference, under the assumption that the quality of the chain is constant across locations. Taking logs of Equation (5) and differencing with respect to the reference establishment, we obtain

$$\log s_{k|sj} - \log s_{0|sj} = \frac{1}{\rho} (\delta_{sjk} - \delta_{s0}) \quad (24)$$

From here, we identify the establishment fixed effects δ_{sjk} for all s , j , and k , up to scale (with the scale governed by the nesting parameter ρ). Once we identify the scaled establishment fixed effects, we can recover the sector-location unobservables ξ_{sj} for all s and j using Equations (4) and (21).

Note that absent additional data or assumptions, δ_{sjk} and ρ are not separately identified, since in practice, we do not observe location and establishment choices separately. In principle, with multi-period data, changes in the set of establishments within a sector-location

could provide identifying variation in within-nest substitution and thus pin down ρ .

5.1.3 Estimation Simplifications

For the current version of demand estimation, we work with a simplified implementation of the model. First, we group establishments into only two broad sectors, restaurants and retail shops. Second, we restrict the estimation sample to consumers for whom both home and work locations can be inferred, so that trip origins are observed on both dimensions. Third, we estimate establishment fixed effects only in CBGs with at least one chain establishment, which serves as a reference point for identifying relative quality within the location. Finally, we assume $\rho = 1$. These simplifications allow us to estimate a transparent baseline version of the model while preserving the main sources of variation relevant for demand estimation.

5.2 Supply

The supply-side parameters are $\Theta = (\mu, \sigma_\delta, \sigma_\nu, \alpha, \beta, \gamma_0, \gamma_1, \gamma_2, \gamma_3)$. We estimate them in two stages. In Stage 1, we use maximum likelihood to recover the entry thresholds $\{\tau_j^*\}$ and the quality distribution parameters $(\mu, \sigma_\delta, \sigma_\nu)$ from the observed distribution of entrant qualities and entry counts. In Stage 2, we use simulated GMM to recover the structural profit parameters $(\alpha, \beta, \gamma_0, \gamma_1, \gamma_2)$ from the zero-profit moment conditions, given some fixed γ_3 .

5.2.1 Stage 1: Identification of μ , σ_δ , σ_ν , and τ_j^*

Entry rate and likelihood. The entry rate at location j satisfies

$$\frac{N_j}{N_j^{\text{pot}}} \xrightarrow{p} p_j^* \equiv \Pr(\tilde{\delta}_{jk} \geq \tau_j^*) = 1 - \Phi\left(\frac{\tau_j^* - \mu}{\tilde{\sigma}}\right)$$

as $N_j^{\text{pot}} \rightarrow \infty$, where $\tilde{\sigma} = \sqrt{\sigma_\delta^2 + \sigma_\nu^2}$. This pins down the standardized threshold $c_j^* \equiv (\tau_j^* - \mu)/\tilde{\sigma}$, which we estimate by $\hat{c}_j = \Phi^{-1}(1 - N_j/N_j^{\text{pot}})$. Since $N_j \sim \text{Binomial}(N_j^{\text{pot}}, p_j^*)$,

by the delta method:

$$\hat{c}_j - c_j^* \approx -\frac{N_j/N_j^{\text{pot}} - p_j^*}{\phi(c_j^*)} \sim \left(0, \frac{p_j^*(1-p_j^*)}{N_j^{\text{pot}} \cdot \phi(c_j^*)^2}\right)$$

so $\hat{c}_j \rightarrow c_j^*$ only as $N_j^{\text{pot}} \rightarrow \infty$, with larger noise for locations with few potential entrants or extreme entry rates.

Substituting $\tau_j^* = \mu + c_j^* \tilde{\sigma}$, the likelihood of entrant quality δ at location j is:

$$\log f(\delta \mid j; \Theta_1) = \underbrace{-\log \sigma_\delta - \frac{(\delta - \mu)^2}{2\sigma_\delta^2}}_{\log \Pr(\delta)} + \underbrace{\log \left(1 - \Phi\left(\frac{\mu + c_j^* \tilde{\sigma} - \delta}{\sigma_\nu}\right)\right)}_{\log \Pr(\nu_{jk} \geq \tau_j^* - \delta)} \quad (25)$$

where $\Theta_1 \equiv (\mu, \sigma_\delta, \sigma_\nu)$. The first term is the log density of the quality draw; the second is the log probability that the entry shock ν_{jk} is large enough to induce entry given quality δ .

The full likelihood is:

$$\log \mathcal{L}(\Theta_1) = \sum_j \sum_{k=1}^{N_j} \log f(\delta_{jk} \mid j; \Theta_1) \quad (26)$$

Global identification. We establish that $(\mu, \sigma_\delta, \sigma_\nu)$ are globally identified from the likelihood evaluated at the true c_j^* .

Proposition 2 (Global Identification: Stage 1). *Suppose (C1) there exist at least two locations with $c_j^* \neq c_{j'}^*$, and (C2) δ_{jk} has support $(-\infty, \infty)$. Then $f(\delta \mid j; \Theta_1) = f(\delta \mid j; \Theta'_1)$ for all j and all δ implies $\Theta_1 = \Theta'_1$. Consequently, the expected likelihood is uniquely maximized at data-generating parameters Θ_1^* .*

Proof. Suppose the densities are equal for all j and δ . Consider location j :

$$\begin{aligned} -\log \sigma_\delta - \frac{(\delta - \mu)^2}{2\sigma_\delta^2} + \log \left(1 - \Phi\left(\frac{\mu + c_j^* \tilde{\sigma} - \delta}{\sigma_\nu}\right)\right) \\ = -\log \sigma'_\delta - \frac{(\delta - \mu')^2}{2\sigma_\delta'^2} + \log \left(1 - \Phi\left(\frac{\mu' + c_j^* \tilde{\sigma}' - \delta}{\sigma'_\nu}\right)\right) \end{aligned} \quad (27)$$

Rearranging, as $\delta \rightarrow -\infty$, the left side goes to $\log(\sigma_\delta/\sigma'_\delta)$. Meanwhile, the right side is

not bounded unless δ^2 and δ both have coefficients equal to 0. So for the equation to hold at all δ , it must be that $\sigma_\delta = \sigma'_\delta$ and $\mu = \mu'$.

Now we only need to show $\sigma_\nu = \sigma'_\nu$. Taking the difference of log densities across two locations j, j' with $c_j^* \neq c_{j'}^*$, cancels the quality density terms, which are location-invariant, yielding a condition that, after differentiating with respect to δ and taking limits, implies $\tilde{\sigma}/\sigma_\nu^2 = \tilde{\sigma}'/\sigma_\nu'^2$. Combined with $\sigma_\delta = \sigma'_\delta$ and hence $\tilde{\sigma} = \tilde{\sigma}'$, this gives $\sigma_\nu = \sigma'_\nu$. Global identification then follows from strict positivity of KL divergence. $\square$

Remark 2 (Consistency with estimated $\hat{c}_j$). *proposition 2* establishes identification at the true c_j^* . In practice we substitute $\hat{c}_j$, which introduces a difference in the log likelihood at each observation:

$$\log \hat{f} - \log f \approx -h(z_{jk}^*) \cdot \frac{(\hat{c}_j - c_j^*)\tilde{\sigma}}{\sigma_\nu}$$

where $z_{jk}^* = (\tau_j^* - \delta_{jk})/\sigma_\nu$ and $h(z) = \phi(z)/(1 - \Phi(z))$ is the Mills hazard. Although $\hat{c}_j - c_j^*$ has mean zero, it introduces a nonzero bias in the score through its covariance with $h(z_{jk}^*)$: since z_{jk}^* depends on c_j^* , which is correlated with $\hat{c}_j$, the errors do not cancel even in expectation. This bias is of order $O(1/N_j^{pot})$ and vanishes as $N_j^{pot} \rightarrow \infty$. The MLE $\hat{\Theta}_1$ therefore remains consistent under joint asymptotics $J \rightarrow \infty$ and $\min_j N_j^{pot} \rightarrow \infty$, but is biased in finite samples, with larger distortions for locations with small N_j^{pot} or extreme entry rates.

Weak identification of σ_ν . Although σ_ν is globally identified in the population, it could be weakly identified in finite samples. The Fisher information with respect to σ_ν is:

$$\mathcal{I}_{\sigma_\nu \sigma_\nu} = \mathbb{E} \left[\left(\sum_{jk} h_{jk} \cdot \frac{\mu - \delta_{jk}}{\sigma_\nu^2} \right)^2 \right] \quad (28)$$

where $h_{jk} = \phi(z_{jk})/(1 - \Phi(z_{jk}))$ and $z_{jk} = (\tau_j^* - \delta_{jk})/\sigma_\nu$. As $\sigma_\nu \rightarrow \infty$, $z_{jk} \rightarrow c_j^*$ (a location-level constant), so $h_{jk} \rightarrow h(c_j^*)$ no longer varies with δ_{jk} , and $\sum_k h_{jk}(\mu - \delta_{jk}) \rightarrow h(c_j^*) \sum_k (\mu - \delta_{jk}) \approx 0$. Hence $\mathcal{I}_{\sigma_\nu \sigma_\nu} \rightarrow 0$: the likelihood becomes flat in σ_ν for large values, and the MLE is unreliable when the true σ_ν is large relative to $\sigma_\delta/|c_j^*|$.

5.2.2 Stage 1: Maximum Likelihood Estimation

Likelihood. We maximize the log likelihood jointly over $(\tau_j)_{j: N_j > 0}$, σ_δ , σ_ν , and μ using a gradient-based method (L-BFGS-B). The full likelihood combines two sources of information:

$$\log \mathcal{L} = \underbrace{\sum_j (N_j^{\text{pot}} - N_j) \log \Phi\left(\frac{\tau_j - \mu}{\tilde{\sigma}}\right)}_{\text{entry counts}} + \underbrace{\sum_j \sum_{k=1}^{N_j} \left[\log \phi\left(\frac{\delta_{jk} - \mu}{\sigma_\delta}\right) + \log\left(1 - \Phi\left(\frac{\tau_j - \delta_{jk}}{\sigma_\nu}\right)\right) \right]}_{\text{entrant quality distribution}} \quad (29)$$

The two terms mirror the two identification steps. The entry count term pins down the standardized threshold $c_j^* = (\tau_j^* - \mu)/\tilde{\sigma}$ at each location, via the observed entry rate N_j/N_j^{pot} . The entrant quality term then separately identifies $(\mu, \sigma_\delta, \sigma_\nu)$ from the shape of the entrants' quality distribution as established in [proposition 2](#). Intuitively, μ shifts the location of the distribution, the spread of δ_{jk} in the right tail where it is unlikely that noise affects the entry pins down σ_δ , while the smoothness of the entry probability around τ_j^* pins down σ_ν .

Initialization. We initialize τ_j at the median of observed δ_{jk} within each location, σ_δ at the sample standard deviation of all observed δ_{jk} , σ_ν at $0.5\sigma_\delta$, and μ at the sample mean of observed δ_{jk} . Bounds are set as $\tau_j \in [\min_k \delta_{jk}, \max_k \delta_{jk}]$, $\log \sigma_\delta, \log \sigma_\nu \in [-3, 3]$ (corresponding to $\sigma \in [0.05, 20]$), and μ within five standard deviations of the sample mean.

5.2.3 Stage 2: Identification of α , β , γ_0 , γ_1 , γ_2

Moment conditions. Given $\{\tau_j^*\}$ recovered from Stage 1, the zero-profit condition ([Equation \(14\)](#)) provides J moment conditions:

$$\mathbb{E}[\pi_j(\Theta, \tau^*)] = 0 \quad \forall j \in \mathcal{J} \quad (30)$$

With $J \gg 5$ parameters, the system is over-identified.

Local identification. Let $D = \partial \mathbb{E}[m_j] / \partial \theta|_{\theta^*} \in \mathbb{R}^{J \times 5}$ be the Jacobian of moments with respect to structural parameters. The parameters are locally identified if $\text{rank}(D) = 5$. The profit moments are linear in β and the γ 's:

$$m_j = \beta \mathbb{E}[\log v_j \mid \alpha, \tau^*] - \gamma_0 - \gamma_1 \log w_j - \gamma_2 \log r_j - \gamma_3 \mathbb{E}[\log N_j \mid \tau^*]$$

so these four parameters are identified from cross-location variation in $(\mathbb{E}[\log v_j \mid \alpha, \tau^*], \log w_j, \log r_j)$, provided the implied regressor matrix has rank 4. The agglomeration parameter α enters nonlinearly through how it shifts $\xi_j = \tilde{\delta}_j^{\text{loc}} + \alpha N_j$ and changes the spatial pattern of visits: α is identified if $\partial \mathbb{E}[\log v_j] / \partial \alpha$ is not spanned by the other regressors.

Plug-in bias from first-stage estimation. An empirical challenge is that Stage 2 uses $\hat{\tau}$ (the Stage 1 MLE estimate) rather than the true τ^* . Since $\hat{\tau}_j \neq \tau_j^*$ in finite samples, the moment conditions are evaluated at the wrong thresholds, introducing a systematic bias:

$$\mathbb{E}[m(\theta^*, \hat{\tau})] = \mathbb{E}[m(\theta^*, \tau^*)] + D_\tau \varepsilon_\tau + O(\|\varepsilon_\tau\|^2) = D_\tau \varepsilon_\tau \neq 0 \quad (31)$$

where $\varepsilon_\tau = \hat{\tau} - \tau^*$ is the Stage 1 estimation error and $D_\tau = \partial \mathbb{E}[m] / \partial \tau^*$. The Stage 2 GMM then minimizes $\|D_\tau \varepsilon_\tau + D_\theta(\hat{\theta} - \theta^*)\|^2$, yielding:

$$\text{bias}(\hat{\theta}) \approx -(D_\theta^\top D_\theta)^{-1} D_\theta^\top D_\tau \varepsilon_\tau \quad (32)$$

The magnitude of this bias is amplified by poor conditioning of $D_\theta^\top D_\theta$ (e.g., if regressors w_j and r_j are highly correlated) and by the alignment of $D_\tau \varepsilon_\tau$ with the column space of D_θ .

5.2.4 Stage 2: Simulated GMM

Moment simulation. For a given parameter vector θ and threshold vector τ , we simulate S draws of the city: each draw consists of (i) firms' qualities, noises, and entry decisions under the threshold strategy, (ii) consumer visits to the marginal firm at each location given

realized entry. Then we compute the expected profits:

$$\hat{m}_j(\theta, \tau) = \frac{1}{S} \sum_{s=1}^S \pi_j^{(s)}(\theta, \tau) \quad (33)$$

where $\pi_j^{(s)}$ is the profit of the marginal firm in simulation s , with visits computed via Equation (11) using the simulated entry configuration. As $S \rightarrow \infty$, $\hat{m}_j \rightarrow \mathbb{E}[\pi_j(\theta, \tau^*)]$ by the law of large numbers. The simulation is implemented using JAX for GPU acceleration, enabling up to $S = 20,000$ draws per evaluation.

Objective function. We minimize the sum of squared moments:

$$Q(\theta) = \frac{1}{J} \sum_j \hat{m}_j(\theta, \hat{\tau})^2 \quad (34)$$

Standard errors. We compute standard errors using a sandwich estimator. Let $D = \partial \hat{m} / \partial \theta|_{\hat{\theta}}$ be the numerical Jacobian, evaluated by finite differences at $\hat{\theta}$ with $S = 5,000$ simulation draws. Let $S_j = \hat{m}_j^2$ approximate the variance of the moment at location j . The asymptotic variance of $\hat{\theta}$ is:

$$\hat{V} = (D^\top D)^{-1} D^\top \text{diag}(S_j) D (D^\top D)^{-1} / J \quad (35)$$

and standard errors are $\hat{s}e_k = \sqrt{\hat{V}_{kk}}$.

Calibrated parameters. The logit scale parameter $\rho = 0.8$, consumer travel cost parameters $\beta_{\text{work}} = \beta_{\text{home}} = 1$, and the inverse supply elasticity of real estate $\gamma_3 = 1.5$ are calibrated rather than estimated.

6 Estimation Results

6.1 Demand Results

Travel costs and trip-chaining. [Table 7](#) reports the estimated travel cost parameters. The estimates show that consumers exhibit spatial frictions in both destination choice and trip chaining. For first visits, consumers are significantly more likely to visit establishments located in the same CBG as their home or workplace. Distances from both home and work reduce the likelihood of visiting a location, although the positive quadratic terms suggest that the marginal disutility of distance attenuates at longer distances. Similar patterns arise for trip-chaining. Consumers are significantly more likely to choose a next destination located in the same CBG as their previous stop, indicating local persistence in trip sequences. Distance from the previous destination reduces the probability of choosing a location, again with a convex relationship reflecting diminishing marginal travel costs at longer distances. Overall, these estimates highlight the central role of spatial frictions and proximity in shaping consumer mobility and multi-stop travel behavior.

We estimate $\hat{\beta}_{TC} = 0.06$, implying that trip-chaining matters for destination choice but with a more modest magnitude than in some previous work.¹⁵ One possible reason is that our trip-chaining parameter is identified from a different notion of sequential visits than in previous studies. We define sequential visits using a three-hour window, whereas previous work often defines trips using home- or work-based travel segments that may span a much longer portion of the day. In addition, our model is estimated at a relatively fine spatial scale and includes rich travel-cost controls and fixed effects. As a result, part of the persistence in sequential visits may be absorbed by these components rather than by the trip-chaining parameter β_{TC} itself.

¹⁵For example, Oh and Seo ([2023](#)) reports an estimate of roughly 0.4.

Table 7: Estimation of Travel Costs

	Destination Choice		Trip-Chaining Choice	
	Restaurants	Shops	Restaurants	Shops
	(1)	(2)	(3)	(4)
Same home CBG	0.315*** (0.016)	0.251*** (0.020)		
Distance from home	−0.073*** (0.004)	−0.046*** (0.005)		
Distance from home squared	0.006*** (0.0004)	0.003*** (0.0006)		
Same work CBG	1.340*** (0.011)	1.085*** (0.015)		
Distance from work	−0.049*** (0.004)	−0.037*** (0.005)		
Distance from work squared	0.004*** (0.0004)	0.003*** (0.0006)		
Same previous CBG			2.085*** (0.023)	1.748*** (0.022)
Distance from previous destination			−0.447*** (0.004)	−0.328*** (0.004)
Distance from previous destination squared			0.032*** (0.0004)	0.023*** (0.0005)
Location FE	Yes	Yes	Yes	Yes
Origin (home-work) FE	Yes	Yes	No	No
Previous location FE	No	No	Yes	Yes

Notes: This table reports the estimated travel cost parameters from our demand model. Columns (1) and (2) report parameters governing destination choice for restaurant and shop visits, respectively. Columns (3) and (4) report parameters governing the choice of the next destination in the restaurant and shop sectors, respectively. The reported trip-chaining parameters are scaled by β_{TC} . Robust standard errors

Sector-location unobservables. Figure 4 reports the estimated sector-location unobservables ξ_{sj} for restaurants and shops. A clear spatial pattern emerges: locations in the eastern and northern parts of San Francisco tend to have higher values of ξ_{sj} , particularly for restaurants. This pattern is consistent with the concentration of commercial corridors and consumption amenities in these areas, and it reflects both the desirability of underlying location fundamentals ($\tilde{\xi}_{sj}$) and endogenous supply-side spillovers. In contrast, lower values are more common in primarily residential or less popular areas. Overall, the spatial distribution of ξ_{sj} aligns with the density of consumption amenities in the city and suggests that location-specific unobservables contribute meaningfully to where consumers choose to visit, in addition to observed travel characteristics and spatial frictions.




Figure 4: Estimated Sector-Location Unobservables

Notes: This figure maps the estimated sector-location unobservables (ξ_{sj}) across CBGs in San Francisco for restaurants (Panel a) and shops (Panel b). We report the qualitative values: low, medium, and high. Areas with missing values correspond to locations with insufficient data for estimation.

Establishment quality distribution. We next report the distribution of estimated establishment fixed effects, interpreted as establishment quality, for restaurants and shops in Figure 5. The distributions are dispersed in both sectors, pointing to substantial heterogeneity in establishment appeal. These estimates, however, should be interpreted with some

caution, as they are based on a restricted sample. In the current implementation, establishment fixed effects are identified only for sector-location pairs that contain at least one chain establishment, which serves as a reference for relative establishment quality within the sector-location. This restriction raises a potential selection concern if branded establishments disproportionately locate in some specific areas. In our sample, locations with chain establishments account for roughly 25% of restaurant CBGs and 36% of shop CBGs. An important next step is therefore to develop an approach that identifies establishment quality in sector-locations without branded establishments, which would allow us to broaden the estimation sample and mitigate the selection concern.




Figure 5: Distribution of Establishment Quality

Notes: This figure reports the distribution of estimated establishment fixed effects, or establishment quality, separately for restaurants (Panel a) and shops (Panel b). The dashed vertical line indicates zero, corresponding to the reference establishment (the branded establishment within each sector-location).

Finally, we examine how the dispersion of establishment quality varies with location quality. [Figure 6](#) plots the standard deviation of estimated establishment quality within each sector-location against the estimated sector-location unobservables. In both sectors, locations with higher values tend to exhibit greater dispersion in establishment quality. This is consistent with the idea that more attractive locations support wider ranges of establishments, resulting in greater within-location heterogeneity in performance, which is key to our

identification strategy.




Figure 6: Firm Quality Dispersion and Sector-Location Unobservables

Notes: This figure plots the standard deviation of establishment quality δ_{sjk} within each sector-location against the estimated sector-location unobservable ξ_{sj} . Each point represents a location within the given sector. The fitted lines show positive relationships between sector-location unobservables and the within-location dispersion of establishment quality in both restaurant and shop sectors.

Reduced-form productivity spillovers. We present a reduced-form regression based on Equation (9). As discussed above, in the structural model, sector-location unobservables reflect both underlying location fundamentals and productivity spillovers. In our reduced-form estimation, we proxy for location fundamentals using the within-location dispersion of firm fixed effects. We further assume a simple parametric form for the agglomeration function, namely $A(\delta_j; \alpha_A) = \alpha_N \times N_{sj}$. Table 8 reports the results. When the proxy for location fundamentals is omitted, the coefficient on the number of firms is larger. Once the dispersion measure is included, the estimated coefficient on firm counts declines. This suggests an upward bias in spillover estimates when location fundamentals are not properly accounted for.

Table 8: Reduced-Form Estimation of Productivity Spillovers

Sector-location Unobservables	Restaurants		Shops	
	(1)	(2)	(3)	(4)
Number of establishments	0.0268*** (0.0045)	0.0332*** (0.0044)	0.0480*** (0.0059)	0.0535*** (0.0058)
Establishment quality dispersion	0.4242** (0.1991)		0.3142* (0.1620)	
Constant	1.9527*** (0.1899)	2.1843*** (0.1116)	1.1233*** (0.2517)	1.4479*** (0.0991)

Notes: This table reports reduced-form regressions of estimated sector-location unobservables, ξ_{sj} , on the number of firms in the sector-location, N_{sj} , and the within-sector-location standard deviation of firm fixed effects, $SD(\delta_{sjk})$, separately for restaurants and shops. Columns (1) and (3) report specifications that control for both firm counts and within-location firm quality dispersion, while columns (2) and (4) omit the dispersion measure and include only firm counts. Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

6.2 Supply Results: Simulation Validation

We validate the estimator using a simulated city before applying it to the Veraset data.

Setup. The simulated city consists of $J = 144$ locations arranged on a 4×4 grid of neighborhoods, each containing 9 sub-locations. Neighborhoods are classified as CBD (center), urban (interior), or suburban (perimeter). The home-work matrix is generated by a gravity model with population-weighted flows, and location qualities are drawn from neighborhood-type-specific normal distributions. The true parameter values are: $\sigma_\delta = 1.0$, $\sigma_\nu = 2.5$, $\mu = 0$, $\alpha = 0.2$, $\beta = 2.0$, $\gamma_1 = \gamma_2 = 0.3$, $\gamma_3 = 0.2$ (fixed), $\gamma_0 = 3.0$, $\rho = 0.5$, $\beta_{\text{work}} = \beta_{\text{home}} = 0.5$.

Stage 1 recovery. [Table 9](#) reports Stage 1 estimates. σ_δ is well-recovered with a bias of +0.01. σ_ν exhibits a downward bias of -0.17 . The estimated entry thresholds have RMSE = 0.75 and a noise-to-signal ratio of 0.82, reflecting substantial finite-sample uncertainty even with 144 locations. Running 100 Monte Carlo simulations yields similar estimates and small standard errors.

Table 9: Stage 1 Recovery: Simulated Data

Parameter	True	Estimated	Bias
σ_δ	1.00	1.01	+0.01
σ_ν	2.50	2.33	-0.17
τ^* (RMSE)	—	—	0.75
τ^* (noise/signal)	—	—	0.82

Notes: This table reports Stage 1 MLE estimates for the simulated city with $J = 144$ locations. RMSE and noise-to-signal ratio are reported for the entry thresholds τ^* .

Stage 2 recovery. [Table 10](#) compares Stage 2 estimates with the true parameter values, under two scenarios: using the true τ^* (infeasible benchmark) and using the estimated $\hat{\tau}$ (feasible estimator).

Table 10: Stage 2 Recovery: Simulated Data

Parameter	True	True τ^*	Estimated $\hat{\tau}$	Bias ($\hat{\tau}$)
α	0.20	0.21	0.29–1.2	+0.09–+1.0
β	2.00	2.01	1.7–2.3	± 0.3
γ_1	0.30	0.31	0.2–0.5	varies
γ_2	0.30	0.29	0.2–0.5	varies
γ_0	3.00	3.02	2.5–4.1	varies
RMSE	—	0.04	0.13–1.4	—

Notes: This table compares Stage 2 GMM estimates with true parameter values under two scenarios: using the true τ^* (infeasible benchmark) and using the estimated $\hat{\tau}$ (feasible estimator). Ranges for the feasible estimator reflect sensitivity to starting values.

With true τ^* , the GMM recovers all parameters with $\text{RMSE} \approx 0.04$, confirming that the moment conditions are informative given correct thresholds. With estimated $\hat{\tau}$, RMSE rises to 0.13–1.4 depending on the starting point, consistent with the bias formula ([Equation \(32\)](#)): the Stage 1 noise (ε_τ , noise/signal = 0.82) is amplified by the condition number of $D_\theta^\top D_\theta$ (≈ 496 in our simulated city) to produce substantial parameter bias. The rent elasticities γ_0 and γ_1 are individually imprecise but their sum $\gamma_0 + \gamma_1$ is better identified, consistent with the near-collinearity of $\log w_j$ and $\log r_j$.

6.3 Supply Results: Veraset Data

Data construction. We estimate the model separately for restaurants and retail shops using Veraset GPS data for San Francisco in June 2019. Firm-level quality δ_{jk} is recovered from the demand-side gravity regressions, using chain-affiliated establishments as normalization (chain establishments all have quality 0). We restrict to CBGs with at least

one chain establishment in the relevant sector, as these provide the reference quality necessary to identify relative establishment quality; this covers approximately 25% of restaurant CBGs and 36% of shop CBGs. The number of potential entrants at each location is set to $N_j^{\text{pot}} = \lceil K \times \max_{j' \in \mathcal{N}(j)} N_{j'} \rceil$, where the maximum is taken over CBGs in the same neighborhood as j and K is some scale parameter. We present results for restaurants as the baseline sector and report sensitivity analyses with respect to K .

Stage 1 results. Table 11 reports Stage 1 MLE estimates for restaurants across different N_j^{pot} scales. The mean quality μ is estimated to be substantially negative across all specifications, ranging from -0.96 (scale $1.0\times$) to -2.39 (scale $5.0\times$). This reflects selection: observed entrants are drawn from the right tail of the quality distribution, so the unconditional mean of entrant quality exceeds the population mean μ . The estimated σ_δ ranges from 1.42 to 1.79, and σ_ν from 1.59 to 1.99. The entry thresholds $\hat{\tau}_j$ become more negative as the scale increases, reflecting the fact that higher N_j^{pot} implies more non-entrants, requiring lower thresholds to rationalize the observed entry rates.

Table 11: Stage 1 Estimates: Restaurants (Veraset SF)

N_j^{pot} scale	$\hat{\mu}$	$\hat{\sigma}_\delta$	$\hat{\sigma}_\nu$	$\bar{\hat{\tau}}$
1.0 $\times$	-0.962	1.417	1.992	-3.929
1.5 $\times$	-1.464	1.572	1.924	-1.251
2.0 $\times$	-1.838	1.667	1.735	-0.858
3.0 $\times$	-2.380	1.794	1.589	-0.545
5.0 $\times$	-2.395	1.688	1.976	0.511

Notes: This table reports Stage 1 MLE estimates for restaurants across different N_j^{pot} scales. $\bar{\hat{\tau}}$ is the mean of estimated location-level entry thresholds.

Stage 2 results. [Table 12](#) reports Stage 2 GMM estimates for restaurants under different N_j^{pot} scales.

Baseline results ($2.0\times$ scale). This baseline specification yields $\hat{\alpha} = 0.087$, implying that each additional entrant at a location raises the location’s attractiveness (in log-visits terms) by 0.087. This is consistent with the reduced-form estimates of agglomeration elasticities in the literature, which has mean 0.096 and median 0.072 ([Rosenthal and Strange, 2004](#)). The revenue elasticity $\hat{\beta} = 1.001$ implies that a 1% increase in expected visits raises profit by approximately 1 percentage point. The rent elasticities $\hat{\gamma}_1 \approx \hat{\gamma}_2 \approx 0$ suggest that conditional on visits (which already capture much of the location’s desirability), worker and resident density add little additional rent pressure beyond what is already captured by $\gamma_3 \log(N_j + 1)$. The fixed cost $\hat{\gamma}_0 = -0.565$ simply absorbs the levels of the regressors.

Table 12: Stage 2 Estimates: Restaurants (Veraset SF)

Scale	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}_1$	$\hat{\gamma}_2$	$\hat{\gamma}_0$	GMM obj.
1.0×	0.421	0.251	0.027	0.125	−9.549	235.2
	(0.003)	(0.005)	(0.021)	(0.018)	(0.181)	
1.5×	5.645	0.155	0.000	0.096	−7.111	253.9
	(10.034)	(0.170)	(0.019)	(0.019)	(3.145)	
2.0×	0.087	1.001	0.000	0.001	−0.565	719.4
	(0.006)	(0.019)	(0.031)	(0.033)	(0.213)	
3.0×	0.165	0.406	0.000	0.000	−0.223	637.1
	(0.011)	(0.013)	(0.024)	(0.031)	(0.200)	
5.0×	0.001	0.489	0.000	0.014	−0.266	180.7
	(0.004)	(0.010)	(0.018)	(0.016)	(0.135)	
<i>Calibrated: $\rho = 0.8$, $\beta_{\text{work}} = \beta_{\text{home}} = 1$, $\gamma_3 = 1.5$</i>						

Notes: This table reports Stage 2 GMM estimates for restaurants under different N_j^{pot} scales. Standard errors in parentheses, computed via sandwich estimator using $S = 5,000$ simulation draws. GMM obj. is the value of the objective function (Equation (34)) evaluated at the optimum.

Limitations. Several caveats apply to the real-data estimates. First, Stage 1 estimates τ^* in a finite sample, introducing plug-in bias into Stage 2. Second, the N_j^{pot} construction introduces further measurement errors in c_j , which correlate across CBGs sharing the same neighborhood, therefore also biasing all Stage 1 estimates. Third, firms’ qualities treated as data in both stages can also have measurement errors, as they are estimated from observed visits in only one month of data (June 2019) and normalized by chain establishments’ quality being 0.

7 Counterfactual: Work-from-Home Adoption

We use the estimated model to ask how the retail landscape would reorganize if a share of San Francisco’s workforce stopped commuting. WFH relocates workers from their workplace CBG to their home CBG, leaving the resident distribution and the city’s total population unchanged. Entry responds to changes in expected visits, as consumers’ origins are shuffled, and changes in rent, as real estate competition cools down in locations previously dense with workers. Consumers face different trip choices, as the firms that decide to enter are different for each location. Consumer welfare is also affected by changes in distances, partly due to their new work locations and partly due to the new visit choices.

7.1 Design

A counterfactual changes the home–work assignment matrix $\mathbf{HW}$, which determines both where consumers originate trips and, through w_j and r_j , the rent term in the entry cost. Everything else is held at its estimated value: the quality distribution parameters $(\mu_\delta, \sigma_\delta, \sigma_\nu)$, the profit parameters $(\alpha, \beta, \gamma_0, \gamma_1, \gamma_2)$, the location fundamentals $Q_j = \tilde{\delta}_j^{\text{loc}} - \alpha \tilde{N}_j$, and the distance matrix.

We do not hold the entry threshold τ_j^* fixed. It is an equilibrium object, the quality cutoff at which a potential entrant’s expected profit equals zero, so in every scenario we re-solve $\mathbb{E}[\pi_j(\tau)] = \xi_j$ to a fixed point.

Two implementation choices deserve comment. First, Stage 2 drives the *average* moment to zero, so the estimated thresholds do not exactly rationalize entry location by location. We therefore treat $\xi_j = \mathbb{E}[\pi_j(\hat{\tau}, \mathbf{HW}_0)]$ as a fixed location-specific cost shifter, standard practice for unobserved costs. The baseline then reproduces observed entry exactly, and scenario differences are attributable to the shock rather than to model fit. Second, the equilibrium is a simulated fixed point carrying Monte Carlo error of order $1/\sqrt{n_{\text{sim}}}$. We pair the simulation seed across baseline and counterfactual so that this error is common and differences out.

We consider three population-preserving shocks. The main scenario, WFH(ϕ), moves a share ϕ of each resident’s work time to their home CBG: $\mathbf{HW} \mapsto (1 - \phi)\mathbf{HW} + \phi \text{diag}(\mathbf{r})$. The residential marginal is untouched and the workplace marginal collapses toward it. We report $\phi \in \{0.25, 0.50, 0.75\}$. Two auxiliary scenarios isolate channels. *Decouple* mixes $\mathbf{HW}$ toward $\mathbf{rw}'/T$, holding *both* marginals fixed so that $\log w_j$ and $\log r_j$ do not move and only the commuting pattern changes. *Suburbanize* resorts residences toward the periphery with each household’s conditional workplace distribution held fixed. Because WFH moves demand and cost together (w_j enters the cost function through $\gamma_1 \log w_j$), we additionally solve demand-only and cost-only variants that hold the other margin at baseline.

7.2 Aggregate Effects

Table 13 reports outcomes across scenarios. The aggregate number of firms barely moves: entry falls by 0.58% at $\phi = 0.25$ and 0.23% at $\phi = 0.50$, then rises by 0.69% at $\phi = 0.75$. All three are small relative to the reallocation underneath them, and we read the aggregate count as close to unchanged rather than as evidence of a meaningful citywide contraction. Visits and consumer surplus, by contrast, rise monotonically in ϕ : total visits increase by 0.38%, 0.76%, and 1.15%, and consumer surplus by 2.68%, 5.36%, and 8.05%, with the surplus gain roughly linear in the WFH share (Figure 7, middle panel). The composition of entrants also shifts down. Visit-weighted entrant quality falls from 1.561 at baseline to 1.505 at $\phi = 0.75$, a decline of 3.6%, as entry becomes less selective in the locations that gain demand.

Consumer surplus here is the expected inclusive value of the nested location–firm choice problem, aggregated over consumers and converted to a percent change. It captures two things: consumers reach firms at lower travel cost, and they face a different local choice set. It does not include the value of the commute itself, which is the largest single component of the welfare effect of remote work in most estimates and which our framework does not model. It also holds residences fixed, so any sorting the shock would induce is excluded, and it excludes landlords’ rental income, which falls where w_j falls. The first exclusion almost

certainly biases our number downward; the third biases it upward. We therefore read the 2.7–8.0% range as the retail-access component of the welfare effect, not as a total.

Table 13: Counterfactual Outcomes by Scenario

Scenario	Δ Entrants (%)	Δ Visits (%)	Δ CS (%)	VW quality
Baseline	0.000	0.000	0.000	1.561
WFH 25%	-0.575	0.380	2.677	1.542
WFH 50%	-0.234	0.763	5.364	1.524
WFH 75%	0.694	1.147	8.048	1.505
WFH 50%, demand only	0.000	0.770	5.449	1.522
WFH 50%, cost only	-0.260	-0.007	-0.036	1.559
WFH 50%, offices empty	10.728	0.787	5.772	1.492
Decouple 50%	0.000	-0.328	-5.391	1.579
Suburbanize 30%	-0.487	-0.194	-2.880	1.559

Entrants, visits, and consumer surplus (CS) are reported as percent changes from baseline. Visit-weighted entrant quality (VW quality) is reported in levels, with the percent change in parentheses in the text. ϕ denotes the share of work time relocated to the residence CBG. *Decouple* holds both the residence and workplace marginals fixed and changes only the commuting pattern. *Demand only* and *cost only* hold the other margin at its baseline value. *Suburbanize* resorts residences toward the periphery. *Offices empty* additionally sets the rent-term worker count to $(1 - \phi)w_j$. All scenarios share a solver seed with the baseline.

7.3 Firm Reallocation

While the citywide firm count remains largely unchanged, our counterfactuals see firms reallocate sharply across locations. In our model, this is driven by entry thresholds changing in response to expected visits and rents.

Figure 7 (left panel) plots the entry response separately for all locations and for the job-dense quartile. The two series behave very differently. In the job-dense quartile, entry rises monotonically in ϕ , reaching +6.5% at $\phi = 0.75$. Across all locations the response is U-shaped: negative at $\phi = 0.25$ and $\phi = 0.50$, positive at $\phi = 0.75$.

The U-shape is the sum of two opposing forces. The demand shock alone would move entry *toward* the residential CBGs that gain workers, so the observed movement in the

opposite direction must run through the rent term. WFH reduces the number of workers in job-dense CBGs, which lowers the rent term $\gamma_1 \log w_j$ there and makes entry cheaper. At the same time it raises the number of workers in residential CBGs, which raises their rent term and makes entry more expensive. Because $dw_j = \phi(r_j - w_j)$, the shock has opposite signs in the two groups: 40 CBGs lose workers and 78 gain them, and the shift is strongly negatively correlated with baseline worker density ($\rho = -0.805$ at every ϕ). At low ϕ the losses in the many gaining-population CBGs dominate the citywide count; by $\phi = 0.75$ the concentrated gains in the job-dense core dominate.

At $\phi = 0.50$, entry in the core rises from 1000.6 to 1036.5 firms while the periphery falls from 1036.8 to 996.2, a reallocation of roughly 3.6% and -3.9% respectively, an order of magnitude larger than the -0.23% aggregate. The grouping is by job density rather than geography, though the two are correlated in San Francisco: [Figure 8](#) shows entry gains falling largely in the northeastern commercial district, while consumer-surplus gains accrue to residents of the southern and western neighborhoods, who no longer travel downtown to consume. Note also that the split is not clean on the population margin: 40 of 118 CBGs lose workers while the job-dense quartile contains roughly 30, so a handful of less job-dense locations lose workers as well.




Figure 7: Dose-Response to Work-from-Home Adoption

Notes: Each panel plots a counterfactual outcome against the WFH share ϕ , holding all estimated parameters fixed and re-solving the entry equilibrium at each point. Left: percent change in the number of entrants, for all locations (blue) and for the job-dense quartile of CBGs (orange). Middle: percent change in aggregate consumer surplus. Right: visit-weighted mean quality of entering firms, in levels. The divergence between the two series in the left panel is the source of the non-monotonic aggregate entry response discussed in [Section 7.3](#).




Figure 8: Spatial Incidence of Work-from-Home Adoption at $\phi = 0.50$

Notes: Top left: change in the number of entrants by CBG, relative to baseline. Top right: change in consumer surplus by CBG of residence. Bottom left: entrants by quintile of baseline worker density w_j , baseline versus WFH 50%. Bottom right: kernel density of entrant quality δ_{jk} , baseline versus WFH 50%; the dashed line marks the baseline mean. Entry gains concentrate in the job-dense northeastern core while surplus gains accrue to residents of the southern and western periphery.

8 Conclusion

This paper studies how spatial demand, local competition, and spillovers shape the distribution of retail firms and their quality within and across urban locations. Using mobile-device GPS data and administrative records for San Francisco, we document two empirical patterns:

firm-level visits vary more within Census block groups than across them, and consumers chain multiple stops into a single trip, so firms generate demand spillovers onto their neighbors. Motivated by these facts, we build and estimate a spatial equilibrium model in which consumers make nested location–firm choices and heterogeneous firms enter under incomplete information, and we use the estimated model to simulate work-from-home adoption.

Our estimates highlight the importance of spatial frictions in shaping urban consumption. All else equal, visits are significantly more likely to occur near home, workplace, and previous trip destinations. We also find that more attractive locations typically support a wider range of firm quality rather than uniformly better firms. This is the variation our identification strategy exploits, since location fundamentals shape not only how many firms enter but also the quality distribution within that location. A reduced-form result on the supply side suggests that failing to account for within-location quality dispersion leads to upward-biased estimates of agglomeration economies.

In the counterfactual, WFH adoption leaves the citywide firm count close to unchanged, but firms reallocate sharply across locations. City-center areas gain entrants despite losing local workers, as lower rents in the worker-dense core offset the loss of local demand there. Aggregate consumer surplus rises by 2.7 to 8.0 percent across adoption levels. Where firm counts fall, the firms that do not enter are those expecting the fewest visits, so consumer welfare falls by less than counts alone would suggest. Taken together, these results suggest that WFH adoption is less damaging to city-center retail than the decline in central business districts documented after COVID-19 might imply.

Several limitations of the current implementation suggest directions for future work. First, we identify firm quality only in locations that contain a branded firm, which restricts the estimation sample and raises a selection concern if branded firms are not randomly distributed across locations. Second, the supply-side parameters are sensitive to the assumed size of the potential-entrant pool; we solve the counterfactual at an agglomeration elasticity of zero, so the reallocation we report should be read as a lower bound on what a specification with

positive scale economies would generate. An independent measure of the potential-entrant pool would sharpen identification on this margin. Finally, extending the model dynamically would allow us to study how the retail landscape adjusts over time following urban shocks or policy interventions, and to separate short-run transition paths from long-run equilibria, a consideration emphasized in the dynamic structural literature ([Keane and Wolpin, 2002](#); [Chan, 2013](#)).

References

- Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015). “The Economics of Density: Evidence from the Berlin Wall”. In: *Econometrica* 83.6, pp. 2127–2189.
- Almagro, M. and T. Domínguez-Iino (2025). “Location Sorting and Endogenous Amenities: Evidence from Amsterdam”. In: *Econometrica* 93.3, pp. 1031–1071.
- Athey, S., D. Blei, R. Donnelly, F. Ruiz, and T. Schmidt (2018). “Estimating Heterogeneous Consumer Preferences for Restaurants and Travel Time Using Mobile Location Data”. In: *AEA Papers and Proceedings* 108, pp. 64–67.
- Cao, Y., J. A. Chevalier, J. Handbury, H. Parsley, and K. R. Williams (2024). *Distributional Impacts of the Changing Retail Landscape*. Cowles Foundation Discussion Papers 2415. Cowles Foundation for Research in Economics, Yale University.
- Cardell, N. S. (1997). “Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity”. In: *Econometric Theory* 13.2, pp. 185–213.
- Chan, M. K. (2013). “A Dynamic Model of Welfare Reform”. In: *Econometrica* 81.3, pp. 941–1001.
- City and County of San Francisco (Feb. 2023). *Roadmap to Downtown San Francisco’s Future*. Office of the Mayor, Fact Sheet. URL: <https://www.sf.gov/file/roadmap-downtown-san-francisco-fact-sheet> (visited on Oct. 10, 2026).
- Cook, C. (2024). “Heterogeneous Preferences for Neighborhood Amenities: Evidence from GPS Data”. In: *The Review of Economics and Statistics*, pp. 1–43.
- Couture, V., J. I. Dingel, A. E. Green, and J. Handbury (2025). *Demographic Preferences and Income Segregation*. NBER Working Papers 33386. National Bureau of Economic Research.

- Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2024). “Income Growth and the Distributional Effects of Urban Spatial Sorting”. In: *Review of Economic Studies* 91.2, pp. 858–898.
- Davis, M. A., A. C. Ghent, and J. Gregory (2024). “The Work-from-Home Technology Boon and Its Consequences”. In: *Review of Economic Studies* 91.6, pp. 3362–3401.
- Holmes, T. J. (2011). “The Diffusion of Wal-Mart and Economies of Density”. In: *Econometrica* 79.1, pp. 253–302.
- Keane, M. P. and K. I. Wolpin (2002). “Estimating Welfare Effects Consistent with Forward-Looking Behavior. Part I: Lessons from a Simulation Exercise”. In: *Journal of Human Resources*, pp. 570–599.
- Liu, S. and Y. Su (2021). “The Impact of the COVID-19 Pandemic on the Demand for Density: Evidence from the US Housing Market”. In: *Economics Letters* 207, p. 110010.
- McFadden, D. (1978). “Modelling the Choice of Residential Location”. In: *Spatial Interaction Theory and Planning Models*. Ed. by A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull. North-Holland, pp. 75–96.
- Miyauchi, Y., K. Nakajima, and S. J. Redding (2025). “The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data”. In: *The Quarterly Journal of Economics* 140.4, pp. 2507–2570.
- Nguyen, H.-A. and T. L. Tran (2025). “The Trade-offs of Curbside Parking: Evidence from Demand-Based Pricing”. Working Paper.
- Oh, R. and J. Seo (2023). “What Causes Agglomeration of Services? Theory and Evidence from Seoul”. Working Paper.
- Qian, F. and Y. Su (2025). “Remote Work and Consumer Cities”. Working Paper.
- Ramani, A. and N. Bloom (2021). *The Donut Effect: How COVID-19 Shapes Real Estate*. SIEPR Policy Brief. SIEPR.

- Rosenthal, S. S. and W. C. Strange (2004). “Evidence on the Nature and Sources of Agglomeration Economies”. In: *Handbook of Regional and Urban Economics*. Ed. by J. V. Henderson and J. F. Thisse. Vol. 4. Elsevier, pp. 2119–2171.
- Rosenthal, S. S., W. C. Strange, and J. A. Urrego (2022). “JUE Insight: Are City Centers Losing Their Appeal? Commercial Real Estate, Urban Spatial Structure, and COVID-19”. In: *Journal of Urban Economics* 127, p. 103381.
- Seim, K. (2006). “An Empirical Model of Firm Entry with Endogenous Product-Type Choices”. In: *The RAND Journal of Economics* 37.3, pp. 619–640.
- U.S. Census Bureau (Sept. 2022). *The Number of People Primarily Working From Home Tripled Between 2019 and 2021*. American Community Survey 1-year estimates. URL: <https://www.census.gov/newsroom/press-releases/2022/people-working-from-home.html> (visited on Sept. 29, 2026).
- Veraset (2022). *Visits*. Dewey Data. <https://doi.org/10.82551/TP7E-VA02>.
- Zhu, T. and V. Singh (2009). “Spatial Competition with Endogenous Location Choices: An Application to Discount Retailing”. In: *Quantitative Marketing and Economics* 7.1, pp. 1–35.

Appendix A Supplemental Figures and Tables




Figure 9: Decomposition of Visit Variance by CBG Visit Quantile (Other Restaurants)

Notes: This figure decomposes the variance of establishment visits to other restaurants into within-CBG and between-CBG components, separately by CBG visit quantile. CBGs are grouped by their overall visit levels, so higher quantiles correspond to more popular locations. The figure is constructed using Veraset data for June 2019.




Figure 10: Decomposition of Visit Variance by CBG Visit Quantile (Shops)

Notes: This figure decomposes the variance of establishment visits to retail shops into within-CBG and between-CBG components, separately by CBG visit quantile. CBGs are grouped by their overall visit levels, so higher quantiles correspond to more popular locations. The figure is constructed using Veraset data for June 2019.




Figure 11: Between-to-Within Visit Variance Ratio by Neighborhood (Other Restaurants)

Notes: This figure shows the ratio of within-CBG to between-CBG variation in visits to other restaurants for each neighborhood. A ratio greater than one indicates that there is relatively more between-location variation than within-location variation. Neighborhoods are shaded according to the magnitude of this ratio. Neighborhoods with substantial missing values are omitted because there is insufficient establishment visit data to compute the ratio. The figure is constructed using Veraset visit data for June 2019.




Figure 12: Between-to-Within Visit Variance Ratio by Neighborhood (Shops)
Notes: This figure shows the ratio of within-CBG to between-CBG variation in visit to retail shops for each neighborhood. A ratio greater than one indicates that there is relatively more between-location variation than within-location variation. Neighborhoods are shaded according to the magnitude of this ratio. Neighborhoods with substantial missing values are omitted because there is insufficient establishment visit data to compute the ratio. The figure is constructed using Veraset visit data for June 2019.

Appendix B Technical Appendix

B.1 Decomposition of Firm-Level Visit Variance

In this section, we prove that

$$\text{Var}(X_k) = \mathbb{E}[\text{Var}(X_k \mid j(k))] + \text{Var}(\mathbb{E}[X_k \mid j(k)])$$

We start with writing down the definition of variance

$$\begin{aligned}
\text{Var}(X_k) &= \mathbb{E}\left[(X_k - \mathbb{E}[X_k])^2\right] \\
&= \mathbb{E}\left[(X_k - \mathbb{E}[X_k|j(k)] + \mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k])^2\right] \\
&= \mathbb{E}\left[(X_k - \mathbb{E}[X_k|j(k)])^2 + (\mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k])^2\right. \\
&\quad \left.+ 2(X_k - \mathbb{E}[X_k|j(k)])(\mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k])\right]
\end{aligned}$$

Taking expectations term by term,

$$\begin{aligned}
\text{Var}(X_k) &= \mathbb{E}\left[(X_k - \mathbb{E}[X_k|j(k)])^2\right] + \mathbb{E}\left[(\mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k])^2\right] \\
&\quad + 2 \mathbb{E}[(X_k - \mathbb{E}[X_k|j(k)])(\mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k])]
\end{aligned}$$

The cross term vanishes because

$$\begin{aligned}
&\mathbb{E}[(X_k - \mathbb{E}[X_k|j(k)])(\mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k])] \\
&= \mathbb{E}[\mathbb{E}[(X_k - \mathbb{E}[X_k|j(k)])(\mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k]) \mid j(k)]] \\
&= \mathbb{E}\left[(\mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k]) \underbrace{\mathbb{E}[X_k - \mathbb{E}[X_k|j(k)] \mid j(k)]}_0\right] = 0
\end{aligned}$$

Therefore,

$$\begin{aligned}
\text{Var}(X_k) &= \mathbb{E}\left[(X_k - \mathbb{E}[X_k|j(k)])^2\right] + \mathbb{E}\left[(\mathbb{E}[X_k|j(k)] - \mathbb{E}[X_k])^2\right] \\
&= \mathbb{E}[\text{Var}(X_k \mid j(k))] + \text{Var}(\mathbb{E}[X_k \mid j(k)]) \quad \square
\end{aligned}$$

B.2 Derivation of the First-Visit Regression Equation

Let N_{hw} denote the total number of consumers with origin (home-work) hw , and let N_{sjhw} denote the number of consumers with origin hw who draw sector s and choose location j . Let $s_{s|hw}$ denote the share of consumers with origin hw who draw sector s , and let $s_{j|shw}$ denote the share of consumers with origin hw who draw sector s and choose location j (with $j = 0$ representing the

outside option).

- For consumers who choose the outside option (i.e., do not visit any location in the city):

$$\begin{aligned} N_{s0hw} &= N_{hw} s_{s|hw} s_{0|shw} \\ &= N_{hw} s_{s|hw} \frac{1}{1 + \sum_{j' \in \mathcal{J}} \exp(\gamma_{shw} + \beta_{\text{home}} \text{Dist}_{hj'} + \beta_{\text{work}} \text{Dist}_{wj'} + \gamma_{sj'})} \end{aligned}$$

- Similarly, for consumers who choose to visit location j :

$$N_{sjhw} = N_{hw} s_{s|hw} \frac{\exp(\gamma_{shw} + \beta_{\text{home}} \text{Dist}_{hj} + \beta_{\text{work}} \text{Dist}_{wj} + \gamma_{sj})}{1 + \sum_{j' \in \mathcal{J}} \exp(\gamma_{shw} + \beta_{\text{home}} \text{Dist}_{hj'} + \beta_{\text{work}} \text{Dist}_{wj'} + \gamma_{sj'})}$$

- Taking the difference of logs yields:

$$\begin{aligned} \log(N_{sjhw}) - \log(N_{s0hw}) &= \gamma_{shw} + \beta_{\text{home}} \text{Dist}_{hj} + \beta_{\text{work}} \text{Dist}_{wj} + \gamma_{sj} \\ \Rightarrow \log(N_{sjhw}) &= \underbrace{\gamma_{shw} + \log(N_{s0hw})}_{\text{"updated" } \gamma_{shw}} + \beta_{\text{home}} \text{Dist}_{hj} + \beta_{\text{work}} \text{Dist}_{wj} + \gamma_{sj} \end{aligned}$$

The regression equation ([Equation \(18\)](#)) then follows with an additional regression error term ε_{sjhw} .

B.3 Derivation of the Second-Visit Regression Equation

Let $N_{s'|j}$ denote the total number of consumers drawing sector s' after visiting location j and $s_{j'|js'}$ denote the share that choose j' upon visiting j and then drawing s' for the next trip.

- For consumers who choose not to continue the trip:

$$N_{js'0} = N_{s'|j} s_{0|js'} = N_{s'|j} \frac{1}{1 + \sum_{\tilde{j} \in \mathcal{J}} \exp\left(\frac{\beta_{\text{cont}}}{\beta_{TC}} \text{Dist}_{j\tilde{j}} + \hat{\gamma}_{js'} + \tilde{\gamma}_{s'\tilde{j}}\right)}$$

- Similarly, for consumers who choose to visit location j' :

$$N_{js'j'} = N_{s'|j} \frac{\exp\left(\frac{\beta_{\text{cont}}}{\beta_{TC}} \text{Dist}_{jj'} + \hat{\gamma}_{js'} + \tilde{\gamma}_{s'j'}\right)}{1 + \sum_{\tilde{j} \in \mathcal{J}} \exp\left(\frac{\beta_{\text{cont}}}{\beta_{TC}} \text{Dist}_{j\tilde{j}} + \hat{\gamma}_{js'} + \tilde{\gamma}_{s'\tilde{j}}\right)}$$

- Taking the difference of logs yields:

$$\begin{aligned}\log(N_{js'j'}) - \log(N_{js'0}) &= \frac{\beta_{\text{cont}}}{\beta_{TC}} \text{Dist}_{jj'} + \hat{\gamma}_{js'} + \tilde{\gamma}_{s'j'} \\ \Rightarrow \log(N_{js'j'}) &= \frac{\beta_{\text{cont}}}{\beta_{TC}} \text{Dist}_{jj'} + \underbrace{\hat{\gamma}_{js'} + \log(N_{js'0})}_{\text{“updated” } \hat{\gamma}_{js'}} + \tilde{\gamma}_{s'j'}\end{aligned}$$

The regression equation ([Equation \(19\)](#)) then follows with an additional regression error term $\varepsilon_{js'j'}$.

Appendix C Counterfactual Decomposition and Robustness

C.1 Decomposition and Incidence

We split the shock to see which margin moves entry and which moves welfare. At $\phi = 0.50$, the demand-only variant leaves aggregate entry unchanged (+0.00%) but delivers almost the entire surplus gain (+5.45%, against +5.36% for the joint shock). The cost-only variant moves entry by -0.26% and surplus by -0.04% . The joint effect on entry, -0.23% , is close to the cost-only effect, so the interaction between the two channels is small.

The welfare gain therefore comes from the demand side: consumers who work from home reach firms from different origins, and surplus rises whether or not firms respond. Entry, in contrast, moves almost entirely through rents. We note that the demand channel alone produces no net entry, so the entry margin contributes less to these results than the surplus numbers alone would suggest.

We use the *decouple* scenario as a placebo. It holds both marginals fixed, so rents and the distribution of trip origins are unchanged and only the commuting pattern moves. Entry does not change, and consumer surplus *falls* by 5.39%: scrambling the home-work pairing lengthens trips without changing where anyone lives or works. The surplus gain under WFH thus comes from shortening the distance between origin and destination, not from reallocating population across the city.

The *offices empty* variant additionally strips the rent-term worker count to $(1 - \phi)w_j$ and produces a 10.7% increase in entry, roughly forty times the baseline WFH 50% response. The gap between the two measures how much of the muted entry response comes from commercial rents not adjusting to the loss of workers. We treat it as an upper bound rather than a prediction, since it assumes rents fall one-for-one with office occupancy.

C.2 Robustness

Simulation error is small. Re-solving the $\phi = 0.50$ equilibrium under three independent solver seeds gives entry changes of -0.27% , -0.26% , and -0.27% and surplus changes of $+5.36\%$ in all three, with a standard deviation below 0.005 on both. Capacity constraints do not bind either: the maximum location fill rate is 0.767 at baseline and 0.808 under WFH 50%, and no location exceeds 90% of its potential-entrant pool in either case.

Parameter uncertainty is more serious. [Table 14](#) re-estimates both stages across values of the potential-entrant scale N^{Pot} , and the Stage 2 parameters are not stably identified. The agglomeration elasticity $\hat{\alpha}$ is 0 at three of five scales and 0.632 at another. The revenue elasticity $\hat{\theta}_{mp}$ ranges from 0 to 0.824. The worker-rent elasticity $\hat{\gamma}_1$ is 0.189 at scale 1.5 and essentially zero at scales 2.0 through 5.0. Stage 1 parameters are more stable ($\hat{\sigma}_\delta$ stays in $[1.42, 1.79]$ and $\hat{\sigma}_\nu$ in $[1.59, 1.99]$), so the instability is specific to the entry-profit parameters on which the counterfactual depends most.

This matters for how we read the results above. We solve the counterfactual at $\hat{\alpha} = 0$, so it contains no agglomeration spillover, and the entry responses we report come from the rent and demand channels alone. We therefore treat them as a lower bound on the reallocation that a specification with positive α would generate. An independent measure of the potential-entrant pool, rather than a scaling of it, would identify the Stage 2 parameters more sharply.

Table 14: Sensitivity of Stage 1 and Stage 2 Estimates to the Potential-Entrant Scale

Scale	$\hat{\mu}$	$\hat{\sigma}_{\delta}$	$\hat{\sigma}_{\nu}$	$\hat{\alpha}$	$\hat{\theta}_{mp}$	$\hat{\gamma}_1$	$\hat{\gamma}_2$
1.0	-0.962	1.417	1.992	0.000	0.108	0.000	0.000
1.5	-1.464	1.572	1.924	0.632	0.000	0.189	0.120
2.0	-1.838	1.667	1.735	0.000	0.706	0.003	0.000
3.0	-2.380	1.794	1.589	0.024	0.824	0.000	0.006
5.0	-2.394	1.688	1.976	0.033	0.733	0.010	0.000

Each row re-estimates both stages with the potential-entrant pool N_j^{pot} scaled by the indicated factor. μ , σ_{δ} , and σ_{ν} are Stage 1 MLE estimates of the quality distribution. α is the agglomeration elasticity, θ_{mp} the revenue elasticity, and γ_1 and γ_2 the worker- and resident-density rent elasticities, all from Stage 2 GMM. Sample is restaurants. The counterfactual in [Section 7.2](#) is solved at $\hat{\alpha} = 0$.