STAFF REPORTS

NO. 1191
APRIL 2026

# Estimating Demand Shocks from Foot Traffic: A Big-Data Approach

Marina Azzimonti | David Wiczer | Yang Xuan

FEDERAL RESERVE BANK of NEW YORK

**Estimating Demand Shocks from Foot Traffic: A Big-Data Approach**
Marina Azzimonti, David Wiczer, and Yang Xuan
*Federal Reserve Bank of New York Staff Reports*, no. 1191
April 2026
https://doi.org/10.59576/sr.1191

**Abstract**

This study leverages high-frequency foot-traffic data from SafeGraph to estimate demand shocks in customer-facing establishments across New York City's retail, service, and health sectors. Recognizing that variations in foot traffic can arise from both unpredictable demand shocks and firm-driven strategies to attract customers, we present a theoretical framework that isolates establishment-level demand fluctuations from firm-level strategic choices. Implementing this empirically, we employ an unsupervised machine learning approach to classify establishments into distinct categories that are largely orthogonal to location and sector. We find important heterogeneity in the persistence of shocks, important heterogeneity in their trends, and that estimation on a pooled sample importantly understates the variance experienced by some establishments.

JEL classification: E21, L14, L80
Key words: consumer-facing brands, service, retail trade, health, demand dynamics, demand shocks, foot traffic

---

Wiczer: Federal Reserve Bank of New York (email: david.wiczer@ny.frb.org). Azzimonti: Federal Reserve Bank of Richmond (email: marina.azzimonti@rich.frb.org). Xuan: Louisiana State University in Shreveport (email: yang.xuan@lsus.edu).

# 1 Introduction

The Service, Retail Trade, and Health sectors are customer-facing industries that account for about one-third of U.S. employment and contribute more value added to GDP than manufacturing.[1] However, demand shocks in these sectors are challenging to measure because output only occurs when a customer arrives at an establishment. Detailed establishment-level data is available for the manufacturing sector, allowing researchers to analyze shocks related to input costs or technology that affect productivity. In contrast, productivity in consumer-facing establishments is driven primarily by the arrival of consumers, which effectively functions as the establishment's productivity process. In other words, a haircut can only occur if a customer with hair arrives at the barbershop, regardless of the availability of combs and skilled hairdressers. To fill this gap, we introduce a novel approach that utilizes high-frequency foot-traffic data to quantify demand shocks within customer-facing sectors across New York City.

We use a newly available mobile-location dataset that associates foot traffic with a particular establishment at a monthly frequency, enabling well-measured estimates of demand at a level of granularity that was not previously feasible. By observing customer arrivals and their characteristics, we uncover the residual component of foot-traffic patterns, which we interpret as establishment-level demand shocks. Because the volatility of these shocks differs across establishments, we use *k*-means clustering to group establishments and estimate shock processes separately for each cluster. Understanding these demand shocks helps establishments adapt strategies, mitigate risks, adjust inputs, and make informed investment decisions. We focus on the period 2018–2019 in order to establish a pre-pandemic baseline for natural variation in demand—something not explored in the existing literature.

Our analysis shows that demand dynamics in customer-facing industries are fundamentally heterogeneous: establishments differ systematically in the persistence, volatility, and growth patterns of their demand processes. These differences cannot be explained by sector, geography, or brand-level strategies. Instead, establishments fall into three economically meaningful demand processes—Fast-Growing, Low-Traffic, and High-Traffic—each implying different volatility profiles and long-run growth prospects. This heterogeneity contrasts to the standard assumption in heterogeneous-firm models that all firms draw shocks from a *single* stochastic process, an assumption that would obscure these differences and misrepresent the underlying distribution of shocks. As models with heterogeneous firms become increasingly dynamic—including applications to investment, firm growth, productivity dispersion, and business-cycle fluctuations—–the dynamic process of shocks faced by such firms becomes increasingly important. The stochastic process governing firm-level shocks plays a central role in determining which establishments are marginal. Our empirical results therefore, provide guidance for calibrating models in which policy responses depend on the distribution of firm-level shocks and on the identity of marginal establishments.

To estimate these processes, we use the *Places* dataset from SafeGraph, which includes detailed foot-traffic data for over 3 million establishments, referred to as "points-of-interest (POIs)," such as retail stores, restaurants, and healthcare facilities, with high spatial and temporal precision. Since SafeGraph tracks the movement of

[1] We use GDP data and employment statistics by sectors from FRED. In this paper, we only consider three customer-facing sectors: Retail Trade (*NAICS 44,45*), Health (*NAICS 62*), and Service (*NAICS 71,72,81*).

individuals over time via cellphone location data, we construct a large-scale panel of store visit patterns. While entirely anonymous, the data provides exceptionally detailed information on the number of visits and number of unique visitors to each establishment, along with business attributes (brand, category, NAICS) and geographic identifiers.

An essential aspect of our study is the focus on *branded establishments*, which allows us to separate random variation in demand from endogenous marketing or product decisions. SafeGraph defines a brand as a "store with multiple locations that share the same logo or store banner." This includes firms that own all their establishments (e.g., Macy's), franchises (e.g., McDonald's), and licensees (e.g., Starbucks). For these firms, product offerings and marketing strategies are typically centrally determined, enabling us to isolate exogenous demand shocks from firm-level strategy.

Our primary technical contribution is to adjust the SafeGraph data to correct sampling biases across locations and time. While SafeGraph provides state-level normalization variables, these are insufficient for smaller geographic markets such as boroughs. By linking the data to the American Community Survey, we construct a time-series census-block-group-level adjustment factor for each borough and additionally create time-specific adjustment factors to correct for biases in time trends.

From there, our empirical analysis is guided by a theoretical framework that separates establishment-level foot-traffic changes attributable to firm-level strategic adjustments, enabling us to estimate demand processes using high-frequency panel data. We examine both year-over-year growth and monthly volatility in the number of visitors for each establishment. A striking observation is the large variation across establishments in within-year growth, especially among high-growth establishments, but also among establishments with nearly zero average growth.

To systematically address this cross-establishment variation, we employ $k$-means clustering. We cluster branded establishments into three groups based on average yearly growth, monthly volatility, and average monthly foot traffic. The first group ("Fast-Growing") exhibits high growth and high volatility. The remaining groups ("High-Traffic" and "Low-Traffic") are distinguished primarily by their traffic levels.

For each cluster we estimate an AR(1) process for year-over-year demand growth, incorporating moving-average error components. High-traffic establishments follow highly persistent, low-variance processes; low-traffic establishments exhibit low persistence and large innovation variance; and fast-growing establishments fall in between. Pooling all establishments overstates both persistence and variance, confounding structural differences in establishment types with stochastic fluctuations.

Importantly, these differences are not explained by borough or sector composition. Each cluster appears across boroughs and sectors in similar proportions. The heterogeneity in demand dynamics therefore reflects establishment-level differences rather than structural differences across industries or locations.

This paper is relevant to researchers who use demand shocks as essential inputs in models of firm dynamics, market structure, and aggregate fluctuations. Our granular, high-frequency measurement of establishment-level demand shocks provides a foundation for calibrating models that depend on realistic representations of heterogeneity in firm-level demand processes.

## 1.1 Literature Review

This paper contributes to the literature seeking to understand the nature and dynamics of demand shocks affecting firms. For instance, Sterk et al. (2021) provide key insights into heterogeneity in the demand process across establishment size. However, while size is often correlated with demand or productivity, it is a less direct measure of the underlying shock, as it reflects firms' endogenous responses to changing demand conditions. Much of the early literature on firm- or establishment-level demand focused on manufacturing, as reviewed in Syverson (2011). Yet, manufacturing accounts for only about 11% of U.S. value added, and the nature of its shocks likely differs from those in retail and service industries, which is the focus of our analysis.

With the advent of new datasets, such as the Longitudinal Business Database (LBD) and ORBIS, researchers have expanded the scope of firm-level revenue analysis beyond manufacturing, as shown by Haltiwanger et al. (2016). However, these sources rely primarily on annual data, which may miss high-frequency changes in demand and are typically released with substantial delays. In contrast, the SafeGraph dataset used in this paper provides near real-time, high-frequency observations, enabling a more granular understanding of demand dynamics.[2] A growing body of research uses location-based data to analyze mobility patterns around establishments, such as Liang et al. (2020). However, these papers remain descriptive, focusing primarily on the arrival patterns. This paper advances that literature by developing a theoretical framework for the demand process and estimating it empirically using high-frequency foot-traffic data.

Our paper is closely related to two recent studies investigating the heterogeneity in firm-level shocks processes. Jaimovich et al. (2023) analyze historical ORBIS data from European countries at an annual frequency, non-parametrically studying firm revenue dynamics and finding significant heterogeneity and variability in shock persistence. Similarly, Melcangi and Sarpietro (2024) study publicly listed U.S. manufacturing firms using Compustat annual sales data and find that smaller firms experience the lowest persistence when their sales are hit by shocks. Complementing these studies, we also identify heterogeneity in shock processes but focus on foot traffic of branded firms in the U.S. within customer-facing sectors rather than revenues. Foot traffic is one step closer to the primitive of a demand shock, capturing the primary driver in customer-facing sectors, while avoiding distortions from pricing, marketing strategies or adjustments to inputs. Moreover, by leveraging granular, location-specific data, our analysis aligns more closely with theoretical models of customer-facing industries.

The process we measure, the arrival of customers at consumer-facing establishments, has been emphasized theoretically but remains empirically underexplored. Specifically, in frictional goods market models, customer arrivals are the key driver of productivity fluctuations, as production in many consumer-facing sectors depends

[2] The COVID-19 pandemic has intensified research using SafeGraph data. Farboodi et al. (2021) use anonymized cellphone data to estimate voluntary social distancing. Goldfarb and Tucker (2020) examine foot-traffic patterns across retail establishments to assess potential COVID-19 transmission risk. Goolsbee and Syverson (2021) investigate the impact of government-imposed restrictions and individual precautionary behavior on economic activity during the COVID-19 pandemic. Cronin and Evans (2020) leverage SafeGraph mobility data to analyze how public restrictions and fear during the COVID-19 pandemic generated heterogeneous demand shocks across retail and service sectors.

on the physical presence of buyers. In these papers, changes in customer arrival serve as a fundamental driving force behind business fluctuations, as shown by Kaplan and Menzio (2016) and Storesletten et al. (2011), amplify other frictions, as in Petrosky-Nadeau and Wasmer (2015), and play a central role in firm dynamics, as demonstrated by Gourio and Rudanko (2014). However, the arrival process itself has remained latent in the data. One of our main contributions is to empirically characterize this process from fundamentals, providing the missing link between theoretical models of customer arrivals and establishment-level evidence.

# 2 Data Description

The establishment-level foot-traffic data used in this study comes from SafeGraph, a private company that aggregates foot-traffic patterns for 5 million businesses and organizations, including 5,500 retail chains and 3 million small businesses.[3] SafeGraph collects GPS location data by continuously pinging roughly 18 million smartphones each day through partnered mobile applications. The resulting dataset tracks movements of individuals between their home Census Block Groups (CBGs) and specific Points of Interest (POIs). In this paper, we use establishment and POI interchangeably. The CBGs are geographical areas that contain a population of about 1,200 individuals. Examples of POIs are retail chains, mom-and-pop businesses, and other public and private establishments, such as religious organizations, schools, and hospitals. Whereas the SafeGraph dataset covers many cities in the United States, the data is most reliable in densely populated areas. Therefore, we restrict attention to three boroughs of New York in our analysis: Manhattan, Brooklyn, and the Bronx.

## 2.1 SafeGraph Data

We use three components of SafeGraph data in this paper. The *Core Places* dataset, which provides establishment-level characteristics for each point of interest (POI), including its location, name, six-digit NAICS code, ZIP code, business category, brand, and a unique POI identifier. These fixed attributes allow us to classify establishments by business type and geographic area. To capture the spatial distribution of consumers, we use the *Home Panel Summary* from the *Neighborhood Patterns* dataset, which includes the information about the number of devices residing in each CBG, the most granular level of the US Census Bureau reports data on. Finally, we use the *Place Patterns* dataset, which reports monthly foot-traffic flows between CBGs and POIs from January 2018 through December 2019. This dataset allows us to trace the origin CBGs of visitors to each establishment over time.

The SafeGraph data are valuable not only because they provide a detailed and high-frequency measure of customer visits to consumer-facing establishments, but also because they can be merged with other public datasets. In this study, we combine SafeGraph with 2019 CBG-level population data from the American Community Survey (ACS) to correct for sampling bias in the raw foot-traffic measures. The adjustment procedure is described in detail in Section 2.2.1.

[3] The data are accessed via Dewey through Advan.

### 2.1.1 An Example: NoMad Starbucks

To illustrate the granularity of SafeGraph data, we examine a specific point of interest (POI): a Starbucks in the NoMad neighborhood of Manhattan. In Figure 1, the dark green lines outline local census block groups (CBGs), our primary geographic units of analysis. The lighter blue boundary marks the CBG containing the Starbucks location.

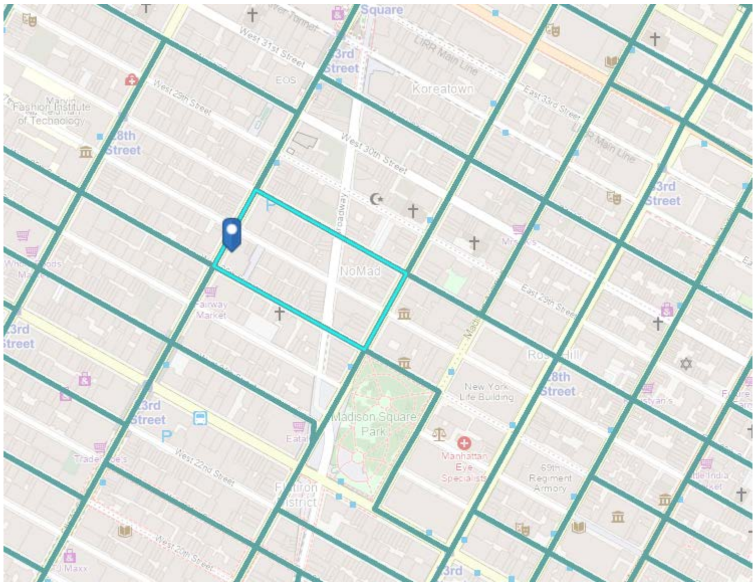


*Note:* Location of a Starbucks in the NoMad area, based on SafeGraph data. The blue square marks its census block group (CBG).

Figure 1: Starbucks Location in NoMad, Manhattan

Throughout the paper, our primary variable of interest is the year-over-year (YoY) change in visitor counts, which helps control for fixed effects and seasonality.[4] Figure 2 plots the year-over-year (YoY) foot-traffic growth for the NoMad Starbucks alongside eight other branded establishments in the Snack and Nonalcoholic Beverage Bars industry (NAICS 722515) located in the same CBG, including several coffee shops near Madison Square Park. The NoMad Starbucks experiences a negative YoY change in foot traffic, which is slightly lower than other similar branded establishments nearby. Its growth is also notably more volatile, suggesting substantial idiosyncratic variation. More broadly, we also find substantial cross-sectional dispersion in foot-traffic trends among similar stores within the same geographic area, with increases and decreases of about 50% not uncommon.

## 2.2 Creating the Branded Sample

This section describes the adjustments made to correct for sampling bias and outlines the sample selection process. We also explain how we identify the branded sample.

[4] To further correct for location-specific and time-varying biases, we construct an adjusted visitor measure, detailed in Section 2.2.1.

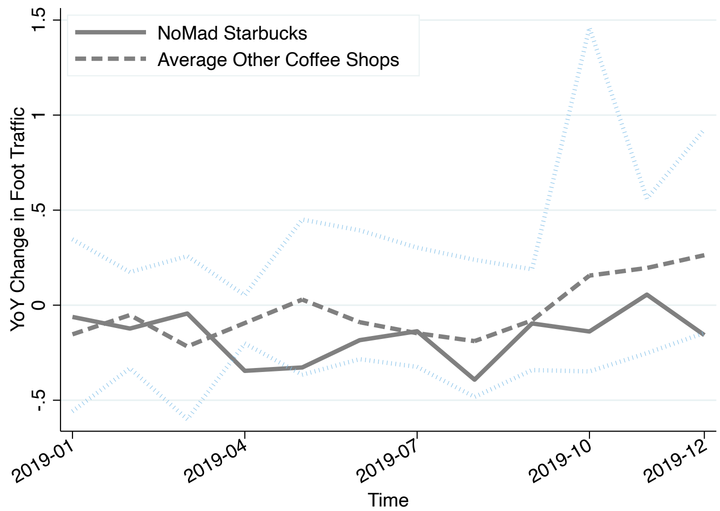


*Note:* Year-over-year (YoY) foot-traffic growth for the NoMad Starbucks compared to other coffee shops in the same CBG, based on SafeGraph data. The gray solid line shows the NoMad Starbucks's trend. The dashed line shows the average, and blue dotted lines indicate the 10th and 90th percentiles of YOY growth rate for other shops.

Figure 2: YoY foot-traffic Growth: NoMad Starbucks vs. Nearby Coffee Shops

### 2.2.1 Sampling Bias

SafeGraph began collecting foot-traffic data in January 2018, progressively expanding coverage by geography, POIs, and the number of tracked mobile devices. This expansion was not representative, introducing two biases: (1) location-specific cross-sectional bias and (2) artificial time-series trends. These distortions appear in the raw foot-traffic data as strong upward trends unrelated to actual economic activity, with patterns varying by borough (Appendix A.1, Figure A12).

SafeGraph provides normalization tools, but only at coarse geographic levels (e.g., state), limiting their usefulness for our granular analysis. We therefore implement our own corrections to address both biases using the American Community Survey (ACS) as a benchmark for representativeness over space and time.

For location-specific bias, SafeGraph assigns each mobile device to an inferred home census block group (CBG). We re-weight CBG-level counts to match ACS population distributions, improving spatial representativeness. For time-series bias, we normalize monthly visitor counts relative to a long-run average, producing a metric that removes artificial upward trends. We refer to this corrected measure as "adjusted toot traffic". All subsequent analysis uses this adjusted series, referred to simply as "foot traffic." Details are provided in Appendix A.2.

### 2.2.2 Sample Selection

We use foot-traffic data from retail, service, and health establishments in Manhattan, Brooklyn, and the Bronx.[5] To ensure data reliability, we exclude establishments with fewer than 10 monthly visitors, extreme fluctuations (e.g., a $\geq$300% monthly increase followed by a reversal), and those in the top 1% of YoY growth. Establishments with fewer than two months of data are also dropped, as growth cannot be computed. Table B4 in Appendix B summarizes the sample selection process and its effect on sample size. Our initial sample includes roughly 1.5 million observations across 71,000 establishments. About 20% of establishments are removed due to data quality filters, and an additional 20% are excluded based on industry selection. After these exclusions, we obtain an unbalanced panel of about 1 million valid establishment-month observations from January 2018 to December 2019, covering over 44,000 unique establishments. We refer to this dataset as the "full sample," which is used in robustness checks.

For our main analysis, we further restrict the sample to establishments with brand information. The term "brand" actually includes a number of corporate structures and franchise arrangements, which we break down in Appendix C. We focus on branded establishments because they operate under centralized strategies—such as standardized pricing, marketing, and product offerings—which help isolate demand shocks from firm-specific endogenous factors. This refinement yields a branded sample of 5,924 establishments and 137,512 observations.[6]

## 2.3 Main Characteristics of the Branded Sample

Table 1 reports summary statistics for adjusted foot traffic per establishment across years, locations, and sectors in the branded sample.

**Establishment Information (POIs).** The first column of Table 1 reports the number of unique POIs. Over half of the branded sample (3,314 of 5,924) are located in Manhattan, 28% in Brooklyn, and the remainder in the Bronx. Service-sector establishments comprise 53% of the sample, retail accounts for 44%, and health establishments just 3%. Restaurants and eating places dominate the sample, consistent with our focus on sectors where production is recorded as revenue only upon an in-person customer visit. Most establishments are corporate-owned or franchised, with limited discretion over pricing and advertising decisions that could endogenously influence foot traffic. Detailed establishment compositions are provided in Appendix C.1.

**Foot Traffic in the Branded Sample.** Table 1 reports average monthly foot traffic and its dispersion across POIs. On average, each branded establishment receives 272

[5] We include establishments with NAICS codes 44-45, 71, 72, 81, and 62. In the service sector, we exclude the *Museums, Historical Sites, and Similar Institutions* (3-digit NAICS code 712), and *Religious, Grantmaking, Civic, Professional, and Similar Organizations* (3-digit NAICS code 813).

[6] An establishment is classified as branded if its brand information is reported in the dataset and the corresponding brand appears in at least two establishments. Branded establishments represent about 8.3% of the full sample but account for nearly a quarter of total foot traffic. Appendix D shows that the branded sample retains many of the key characteristics of the full dataset.

| | # of POIs | Mean FT | St. Dev. FT |
|---|---|---|---|
| Overall | 5,924 | 272 | 381 |
| *By Year* | | | |
| 2018 | 5,924 | 271 | 377 |
| 2019 | 5,924 | 273 | 389 |
| *By Borough* | | | |
| Manhattan | 3,314 | 271 | 413 |
| Brooklyn | 1,679 | 265 | 362 |
| Bronx | 931 | 290 | 282 |
| *By Sector* | | | |
| Retail | 2,608 | 240 | 364 |
| Service | 3,128 | 310 | 398 |
| Health | 188 | 86 | 126 |

*Note:* FT stands for the monthly adjusted foot traffic per establishment.

Table 1: Summary Statistics of Monthly Adjusted Foot Traffic

unique visitors per month, with considerable variation (standard deviation 381). This variation is skewed by a small number of establishments attracting exceptionally high foot traffic, creating a long right tail in the distribution (see Figure C18 in Appendix C.1). Geographically, establishments in Manhattan and the Bronx attract more visitors than those in Brooklyn. As anticipated, establishments in the service sector attract significantly higher visitor counts per store, and the average foot traffic in the health sector is the lowest. Average foot traffic remains stable between 2018 and 2019, both in level and dispersion (Figure C17, Appendix C.1). Although the branded sample is not randomly selected, Appendix D shows that its traffic characteristics closely mirror those of the full sample.

# 3 Towards Demand Shocks: Foot-Traffic Growth

In this section, we then analyze the dynamics of foot-traffic growth, a key indicator of demand shocks. Foot-traffic growth is particularly informative for customer-facing establishments because it directly reflects consumer arrival patterns while avoiding distortions from firm-level factors such as pricing strategies and promotions that affect revenues. Using our branded sample, which provides high spatial and temporal resolution, we examine year-over-year (YoY) variations and monthly volatility in foot traffic across establishments, highlighting its relevance as a foundation for measuring demand dynamics. The year-over-year growth rate of foot traffic for establishment $i$ at time $t$ is defined as:

$$\Delta_{i,t} = \frac{d_{i,t} - d_{i,t-12}}{d_{i,t-12}},$$

where $d_{i,t}$ denotes the monthly foot traffic count for establishment $i$ at time $t$.

The left panel of Figure 3 depicts the distribution of establishments by annual foot-traffic growth from 2018 to 2019 across establishments, which is positively skewed. The quick tapering on the left tail likely reflects sample selection, as the branded establishments in our sample are generally not those experiencing rapid decline.

The right panel of Figure 3 plots the within-year standard deviation of foot traffic conditional on the average growth rate. The nadir corresponds to the modal growth

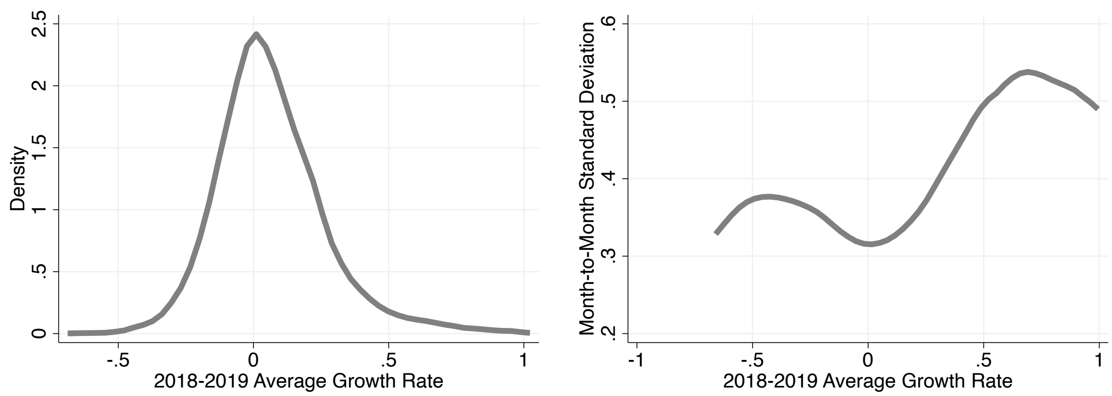


(a) Average YoY Growth

(b) Monthly Volatility

Figure 3: Heterogeneity in Growth Patterns

rate, where foot traffic is relatively stable. However, even at this point, monthly volatility exceeds 30%. Establishments experiencing either growth or shrinking foot traffic exhibit considerably higher volatility than those near the nadir. These findings suggest that changes in demand are far from smooth: even when annual foot traffic nearly doubles, monthly changes remain highly uneven.

## 3.1 Spatial Heterogeneity in Foot-Traffic Growth

Similar to Figure 3, Figure 4 illustrates the distribution of POIs by average annual growth rate (left panel) and the relationship between growth rate and monthly volatility (right panel). The left panel shows positively skewed distributions across all boroughs, with the Bronx exhibiting the lowest kurtosis. Brooklyn has a higher share of establishments with strong year-over-year growth, whereas Manhattan has more non-growing establishments. These patterns are consistent with the summary statistics in Table C.3 of Appendix C.2. The right panel indicates that both shrinking and rapidly growing establishments experience higher monthly volatility in all boroughs. Moreover, Manhattan consistently displays greater volatility than the other boroughs, even for POIs with similar annual growth rates. These findings highlight the heterogeneity in foot-traffic pattern across locations.

## 3.2 Sectoral Heterogeneity in Foot-Traffic Growth

The left panel of Figure 5 shows positively skewed distributions of POIs by average annual growth rates across all sectors. A larger share of POIs in the service sector experience near-zero average growth in foot traffic over two years compared to the other sectors. The right panel reveals that monthly volatility for service-sector establishments is generally lower than for establishments in the other two sectors, regardless of whether foot traffic is growing or shrinking over the two-year period. In contrast, a higher proportion of health-sector establishments exhibit positive average annual growth rates, though their growth patterns display substantial fluctuations. Overall, these patterns indicate considerable heterogeneity in foot-traffic arrival dynamics

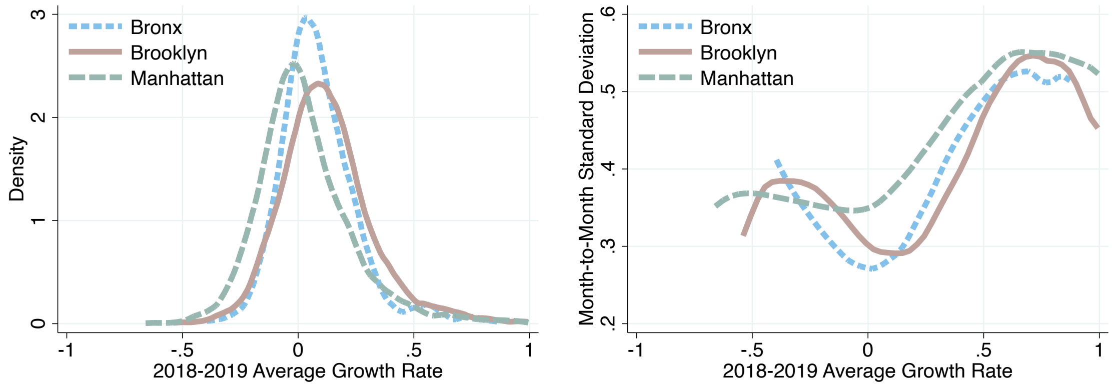


(a) Average YoY Growth

(b) Monthly Volatility

Figure 4: Heterogeneity in Growth Patterns (By Boroughs)

across sectors. We discuss how these dynamics can be inferred and incorporated into demand shock estimation in the next section.

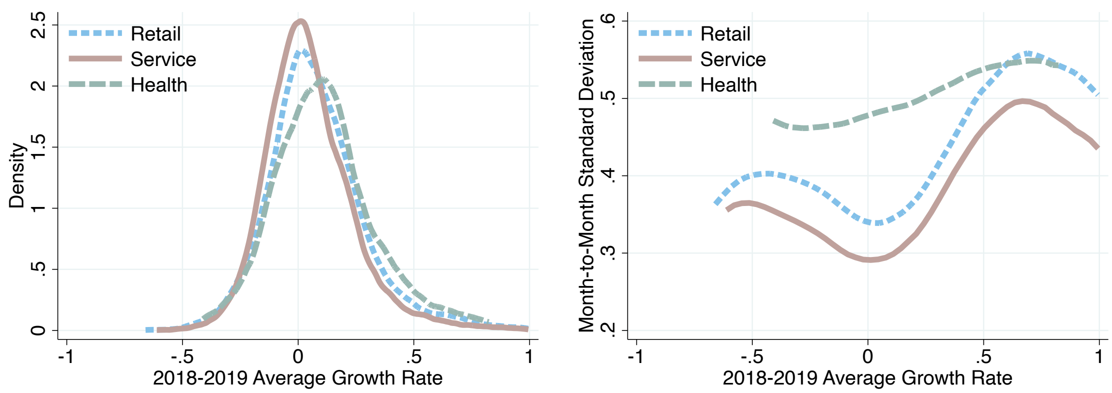


(a) Average YoY Growth

(b) Monthly Volatility

Figure 5: Heterogeneity in Growth Patterns (By Sectors)

# 4 A Model to Facilitate Estimation of Demand Dynamics

The ability to track year-over-year growth in foot traffic across establishments using this newly constructed panel data provides a novel perspective on business dynamics. This granular data reveals substantial monthly volatility in customer visits within individual establishments, highlighting the randomness of consumer behavior. These fluctuations are not mere noise; they stem from a combination of unpredictable demand shocks and firm-level decisions aimed at driving traffic. In customer-facing sectors such as retail, restaurants, and entertainment, demand for goods and services is closely linked to physical customer presence, making foot traffic a valuable proxy for demand dynamics. However, foot traffic is not itself the fundamental process, it reflects both exogenous demand shocks and endogenous firm decisions.

This section establishes a framework to understand the relationship between foot traffic and actual demand for services at the establishment level. This model accounts for firm-level decisions, such as advertising campaigns that can drive increased traffic across all establishments within a firm. Simultaneously, it acknowledges the establishment-specific demand shocks, such as local events, weather patterns, or changes in the competitive landscape, which can cause significant variations in foot traffic at individual locations. This model is an intentionally parsimonious tool for isolating these establishment-level demand shocks from firms' endogenous choices so as to be consistent with other models of service-sector firms for which these shocks can serve as an input.

## 4.1 Firm-Level Advertising Choice Facing A Dynamic Demand Process

Firms are indexed $j$ and establishments indexed $i$, with each establishment having a size denoted by $\mu_i$. The set of establishments belonging to firm $j$ is $\mathcal{J}_j$. These establishments operate in different markets, indexed by $m$. To set concepts, we can use our earlier Starbucks example where the establishment $i$ is the Starbucks in NoMad, $j$ refers to the Starbucks brand and $m$ is the set of other Starbucks locations operating in different markets.

Each firm chooses a level of marketing effort, denoted by $s$, and incurs a cost $\kappa(s)$. After the firm chooses its marketing level, each establishment experiences a random shock, $z_i$, affecting customer arrivals. While the unconditional expectation of these shocks is similar across a given market ($E[z_i] = \bar{z}_m \ \forall i$), they follow a conditional process. Each establishment experiences a heterogeneous trend in log customer arrivals:

$$\log z_{i,t} = \rho_i \log z_{i,t-1} + (1 - \rho_i)\eta_i \times t + \zeta_{i,t},$$

where $\rho_i$ captures the persistence of the shock, $\zeta_{i,t}$ is an establishment-specific innovation.[7] $\eta_i$ captures establishment-level trends in growth, as in those caused by managerial quality. For instance, if a manager is systematically better at utilizing the firm's resources and managing their personnel or if this location is particularly good.[8]

The number of customers arriving at an establishment is a function of marketing effort and the random shock, $d(s_{j,t}, z_{i,t})$. Importantly, establishments cannot sell more than the number of arriving customers, so their output is dictated by $d()$. These advertising efforts are the way in which traffic is endogenous and captures anything from signage to special promotions that draw in consumers. In our Starbucks example, this can be an A-frame sign on the sidewalk given to the establishment by the firm or seasonal drinks and their associated advertising.

Aside from their advertising ability, $s_{j,t}$, firms and establishments are essentially price-takers and they face an infinitely elastic supply of labor and other inputs. This

[7] Notice that the AR(1) assumption is somewhat simpler than the data will want. We could enrich this underlying shock process to be more in-line with the MA specification suggested by the data which will just add terms to the algebra below.

[8] We put all decision making in the hands of the firm below, when they choose advertising. This is a simplification and other franchise-like models might have more establishment effects that we embed in $\eta_i$.

means that in each market, this is a per-unit profit of $\pi_{m,t}$, which is assumed to be positive. These markets are defined quite narrowly here as both the market in which firms sell, and the market in which firms acquire inputs, essentially both a location and a fine industry code.

The firm aims to maximize its expected profit by choosing the optimal marketing level, considering sales, costs, and random shocks. Its optimization problem is:

$$\max_{s} E\left[\sum_{i\in\mathcal{J}_j} d(s,z_{i,t})\pi_{m,t}\mu_i - \kappa(s)\middle| z_{i,t-1}{}_{i\in\mathcal{J}_j}\right].$$

We then specialize the model and assume some functional forms to make our results more implementable. First, foot traffic is multiplicative, $d(s,z) = sz$, and advertising has a quadratic cost $\kappa(s) = \frac{\kappa_0}{2}s^2$. This makes the firm's first-order condition, which sets advertising

$$s = E\left[\sum_{i\in\mathcal{J}_j} z_{i,t}\mu_i \middle| \{z_{i,t-1}\}_{i\in\mathcal{J}_j}\right]\frac{\pi_{m,t}}{\kappa_0}.$$

This means that advertising choice is linear in the firm's expectation of the average establishment shock $z_i$. Putting that back into our foot-traffic function, we get:

$$d(s_{j,t},z_{i,t}) = z_{i,t}E\left[\sum_{k\in\mathcal{J}_j} z_{k,t}\mu_k \middle| \{z_{k,t-1}\}_{i\in\mathcal{J}_j}\right]\frac{\pi_{m,t}}{\kappa_0}.$$

This form presents a multiplicative relationship between the establishment-level shock $z_{i,t}$, firm-level choices, $E\left[\sum_{k\in\mathcal{J}_j} z_{k,t}\mu_k \middle| \{z_{k,t-1}\}_{i\in\mathcal{J}_j}\right]$, and the potential profits they can make, as given to them by the market, $\frac{\pi_{m,t}}{\kappa_0}$.

## 4.2 From Model to Specification

To use our model to inform our view of the data, we convert the first-order condition into an easily implementable regression format. First, we take logs and get

$$\log d(s_{j,t},z_{i,t}) = \log z_{i,t} + \log\left(E\left[\sum_{k\in\mathcal{J}_j} z_{k,t}\mu_k \middle| \{z_{k,t-1}\}_{i\in\mathcal{J}_j}\right]\right) + \log\left(\frac{\pi_{m,t}}{\kappa_0}\right). \quad (1)$$

Define $\hat{d}_{i,t} = \log d(s_{j,t},z_{i,t})$, $\hat{z}_{i,t} = \log z_{i,t}$, $\hat{S}_{j,t} = \log\left(E\left[\sum_{k\in\mathcal{J}_j} z_{k,t}\mu_k \middle| \{z_{k,t-1}\}_{i\in\mathcal{J}_j}\right]\right)$ and the market effects $\hat{v}_{m,t} = \log\frac{\pi_{m,t}}{\kappa_0}$. In general, the market effects, $\hat{v}_{m,t}$, will have periodicity, or seasonality. If we assume that they take the form of

$$\hat{v}_{m,t} = \hat{v}_m + \hat{v}_t + \lambda_{m,t},$$

and $\lambda_{m,t} = \lambda_{m,t-12}$ where $\lambda_{m,t}$ is the seasonal factor and $\hat{v}_t$ is the month-specific innovation and $\hat{v}_m$ the market fixed effect. So here we have assumed to be log-complementary between aggregate time-effects and fixed-market effects.

We now derive the estimating equation from the model. To remove periodicity and fixed market effects, we difference over a 12-month horizon:

$$\Delta_{i,t} = \hat{d}_{i,t} - \hat{d}_{i,t-12} = \hat{z}_{i,t} - \hat{z}_{i,t-12} + \hat{S}_{j,t} - \hat{S}_{j,t-12} + \hat{v}_{m,t} - \hat{v}_{m,t-12}. \tag{2}$$

As a first step, we estimate the following regression to obtain the residual component $\tilde{\Delta}_{i,t}$:

$$\Delta_{i,t} = \log\left(\frac{d_{i,t}}{d_{i,t-12}}\right) = u_i + v_t + \xi_{j,t} + \tilde{\Delta}_{i,t}. \tag{3}$$

Here, $v_t$ denotes the year-over-year change in the market-time fixed effect, $\hat{v}_{m,t} - \hat{v}_{m,t-12} = \hat{v}_t - \hat{v}_{t-12}$.The brand-specific time effect, $\xi_{j,t}$ is the year-over-year difference $\hat{S}_{j,t} - \hat{S}_{j,t-12}$ capturing brand-level advertising decisions. The establishment fixed effect, $u_i = \eta_i \times 12$ reflects trend growth in the $z_{i,t}$ process described above.

Returning to the process for the underlying demand shock, we model year-over-year growth in $\hat{z}_{i,t}$ as:

$$\log\left(\frac{z_{i,t}}{z_{i,t-12}}\right) = \rho_i \log\left(\frac{z_{i,t-1}}{z_{i,t-13}}\right) + (1-\rho_i)\eta_i \times 12 + \epsilon_{i,t}, \tag{4}$$

where the year-over-year innovation is $\epsilon_{i,t} = \zeta_{i,t} - \zeta_{i,t-12}$. We can put this back into Equation (2) and combine with our definition of $\tilde{\Delta}_{i,t}$ from Equation (3) to get our final estimating equation.

$$\Delta_{i,t} = \rho_i \log\left(\frac{z_{i,t-1}}{z_{i,t-13}}\right) + (1-\rho_i)\eta_i \times 12 + \epsilon_{i,t} + \hat{S}_{j,t} - \hat{S}_{j,t-12} + \hat{v}_{m,t} - \hat{v}_{m,t-12},$$

$$\Delta_{i,t} = \rho_i(\Delta_{i,t-1} - v_{t-1} - \xi_{j,t-1}) + (1-\rho_i)u_i + \epsilon_{i,t} + \xi_{j,t} + v_t.$$

Since $\tilde{\Delta}_{i,t} = \Delta_{i,t} - v_t - \xi_{j,t} - u_i$, we have

$$\Delta_{i,t} - v_t - \xi_{j,t} - u_i = \rho_i(\Delta_{i,t-1} - v_{t-1} - \xi_{j,t-1} - u_i) + \epsilon_{i,t}.$$

Substituting the first-stage residual $\tilde{\Delta}$ into this expression yields our estimating equation:

$$\tilde{\Delta}_{i,t} = \rho_i \tilde{\Delta}_{i,t-1} + \epsilon_{i,t}. \tag{5}$$

In this, we have used a first stage to take out market and brand fixed effects and we are left with a dynamic equation for residual foot traffic that can estimate both the persistence $\rho_i$ and innovations $\epsilon_{i,t}$ for the *underlying* demand process. To restate this, this specification now lets us estimate the demand shock process directly from foot traffic growth patterns where $\rho_i$ and $\epsilon_{i,t}$ are actually deep parameters for the demand innovation process that are estimated off of the dynamics of residual foot traffic.

# 5 A Machine Learning Approach to Growth Heterogeneity

We next examine heterogeneity in demand shock process across establishments. Consistent with the literature on establishment-level productivity, which often emphasizes the existence of residual heterogeneity within the manufacturing sector, we anticipate finding residual heterogeneity within the customer-facing sectors.

We use a $k$-means clustering approach to uncover the underlying heterogeneity in demand dynamics. This approach, previously utilized by Bonhomme et al. (2022) to examine wage patterns, allows us to characterize groups of establishments based on similarities in the demand patterns. We choose this unsupervised machine learning approach instead of a parametric approach for two reasons (see Sterk et al. (2021)). Firstly, the data we are working with is novel, so employing an agnostic approach helps us avoid biased modeling choices due to limited prior knowledge. Secondly, given the large cross-section of data available but with a short panel and ongoing improvements in data quality, clustering enables us to identify and isolate outliers without affecting the other parameters.

Following Bonhomme et al. (2022), we first determine the optimal number of clusters $K$, and then estimates the dynamic panel in (5) within each cluster $k$, using residual growth $\tilde{\Delta}_{i,t}$ from (3). This yields cluster-specific persistence estimates $\rho(k)$, which are directly comparable to the structural $\rho_i$ in the model.[9] We form clusters based on three establishment-specific characteristics: average year-over-year growth rate ($\overline{\Delta}_i$), monthly volatility ($\sigma_{\Lambda_i}$), and average foot traffic ($\bar{d}_i$) over two years. The optimal number of clusters, denoted as $K$, is the minimum ensuring that the sum of within-cluster variances is smaller than the sum of within-establishment variances across time. The clustering algorithm identifies 3 distinct categories, and Table 2 outlines the characteristics of the chosen clusters.

We define the group with the highest annual demand growth as the Fast-Growing category. About 25.6% of establishments fall into this cluster, each serving an average of 201 unique customers per month. Within this group, 49.2% are from the retail sector and 47.3% from the service sector, predominantly located in Manhattan and Brooklyn. The Low-Traffic cluster, characterized by the lowest average monthly foot traffic—of about 61 customers per store—, includes 15.5% of establishments. Conversely, the High-Traffic group, characterized by the highest foot traffic—at approximately 359 unique customers per store monthly—, constitutes 58.8% of the establishments. Interestingly, establishments in the High-Traffic group show nearly zero average year-over-year growth in foot traffic, with a smooth monthly trend. Conversely, those in the Low-Traffic cluster, while not experiencing annual growth, exhibit the highest volatility in monthly foot traffic changes. Note the differences in sectoral compositions between the Low-Traffic and High-Traffic clusters: the Low-Traffic group predominantly comprises establishments from the retail sector, while the High-Traffic cluster is largely made up of establishments from the service sector and has a high fraction of establishments from the health sector. Additional characteristics for each cluster, including detailed sectoral composition, top categories, and brand compositions, are provided in Appendix C.4.

Our findings based on equations (3) and (5) indicate the presence of substantial heterogeneity in the demand process. Figure 6 displays the estimates for each cluster and the pooled regression. We observe considerable disparities between clusters and pooled regression estimates. Specifically, the estimate for persistence for the Low-Traffic group is significantly lower than those from the other two groups and the

[9] We reject the i.i.d. error assumption due to autocorrelation and adopt an MA(4) specification, which passes diagnostic tests for all clusters; see Appendix C.5. As we suggested when describing the model, incorporating this into our underlying demand structure is only a matter of adding a richer process for $z$ and complicating the algebra somewhat.

| *Variable:* — *Cluster:* | Fast-Growing | Low-Traffic | High-Traffic | Pooled |
|---|---|---|---|---|
| *FT Characteristics:* | | | | |
| $\overline{\Delta}_i$ | 0.289 | 0.073 | -0.035 | 0.065 |
| $\sigma_{\Lambda_i}$ | 0.316 | 0.718 | 0.245 | 0.337 |
| $\overline{d}_i$ | 201.282 | 61.186 | 358.862 | 272.277 |
| *Sectoral Compositions:* | | | | |
| % in Retail | 0.492 | 0.534 | 0.393 | 0.440 |
| % in Service | 0.473 | 0.383 | 0.590 | 0.528 |
| *Location Compositions:* | | | | |
| % in Manhattan | 0.421 | 0.681 | 0.587 | 0.559 |
| % in Brooklyn | 0.414 | 0.226 | 0.242 | 0.283 |
| *Regression Estimates:* | | | | |
| $\rho$ | 0.596*** | 0.162 | 0.799*** | 0.710*** |
| | [0.127] | [0.136] | [0.073] | [0.081] |
| $\sigma_{\epsilon_{i,t}}$ | 0.439 | 0.542 | 0.265 | 0.410 |
| *Arellano-Bond Test for:* | | | | |
| $AR(4)P-value$ | 0.640 | 0.416 | 0.098 | 0.364 |
| $AR(5)P-value$ | 0.658 | 0.604 | 0.471 | 0.984 |
| *Fraction of Establishments* | 0.256 | 0.155 | 0.588 | 1.000 |

*Note:* Standard errors are reported in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2: Characteristics & Estimates of Clusters

pooled analysis. Moreover, the Low-Traffic group exhibits the highest standard deviation within establishments, whereas the High-Traffic group shows the lowest, indicating that variance within points of interest (POIs) for year-over-year demand growth, calculated as $\frac{\sigma^2_{\epsilon_{it}}}{1-\rho^2}$, is different across groups. These results highlight that estimates from the pooled regression may not accurately represent the data, and there is also heterogeneity between clusters.

We next examine whether the observed differences arise from variation across boroughs or sectors. Figure 7 shows the within-establishment standard deviation of residuals ($\sigma_\epsilon$) by cluster, disaggregated by sector (panel a) and borough (panel b). Across all sectors—retail (green), services (red), and health (yellow)—"High-Traffic" establishments consistently have the lowest standard deviation of shocks. Similarly, across boroughs—Manhattan (green), Brooklyn (red), and the Bronx (yellow)—the "High-Traffic" group again records the lowest volatility, while the "Low-Traffic" group shows the highest. These patterns suggest that neither sectoral nor geographic differences help explain the heterogeneity in demand shocks.

## 5.1 Robustness to Sample Selection

In this section, we use the full sample to verify the robustness of our findings to the sample exclusion of non-branded establishments. Firstly, we categorize establishments in the full sample into these three clusters, keeping the centroids of each cluster as defined previously in the brand sample. Table 3 lists out the characteristics of these three clusters for all establishments in the full sample. According to Table 3, we find the characteristics of these three clusters are similar to those in the brand sample. In the full sample, 19.4% of establishments are categorized as Fast-Growing, with an

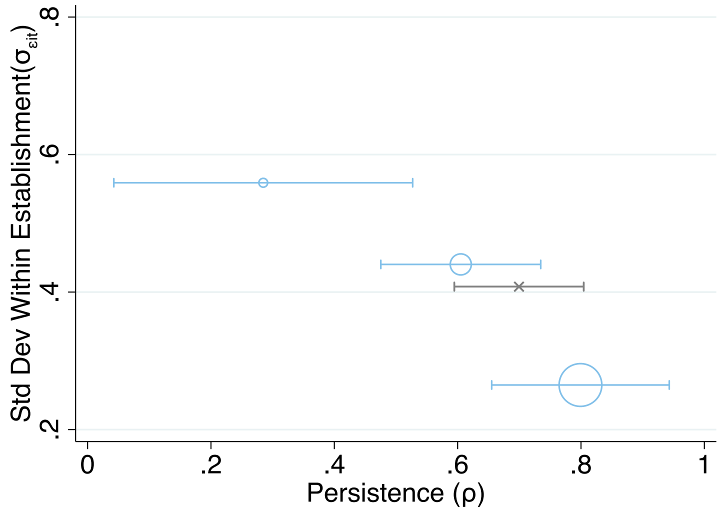


*Note:* The figure illustrates the relationship between the standard deviation within establishments ($\sigma_\epsilon$) and the persistence level ($\rho$). Dot size is proportional to total visitors per cluster (Low-Traffic = smallest, Fast-Growing = mid-size, High-Traffic = largest). Centers indicate cluster-specific estimates; the pooled regression estimate is marked with a bold gray "x." Horizontal bars denote 95% confidence intervals for $\rho$.

Figure 6: Clusters and Pooled Estimates

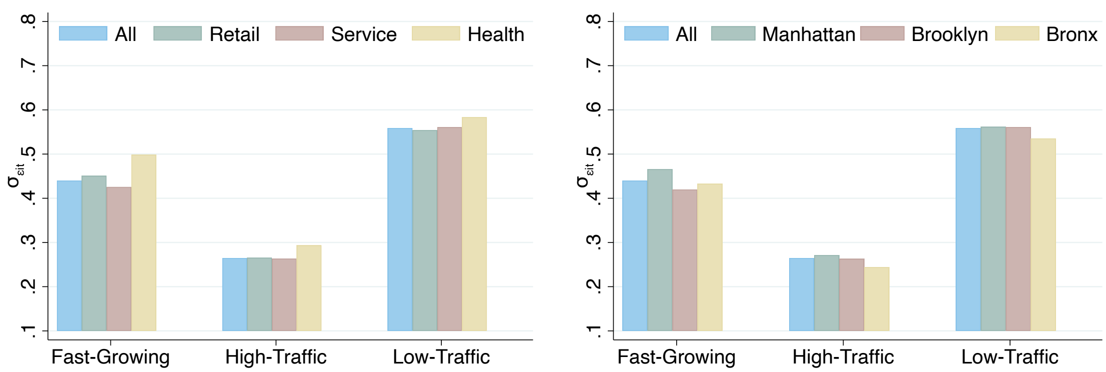


(a) By Sectors (b) By Boroughs

*Note:* The figure shows the standard deviation of residuals for establishments in each cluster, categorized by sectors or boroughs. The horizontal axis in both panels lists cluster groups. Blue bars show the overall cluster-level standard deviation (see Table 2); other bars show the same statistic for each sector or borough within a cluster.

Figure 7: Standard Deviation of Residuals By Groups

average annual growth rate higher than one-third and high monthly volatility. Meanwhile, 22% fall into the Low-Traffic cluster and 58.6% are classified as High-Traffic. Similar to the brand sample, establishments in the High-Traffic group display almost zero average year-over-year growth in foot traffic with a consistent monthly pattern. Conversely, those in the Low-Traffic cluster, despite no annual growth, exhibit the highest volatility in monthly foot-traffic changes. Within each cluster, the distribution of establishments across locations and sectors mirrors that of the brand sample, with the exception of a higher proportion of establishments from the health sector in the full sample.

| *Variable:* — *Cluster:* | Fast-Growing | Low-Traffic | High-Traffic | Pooled |
|---|---|---|---|---|
| *FT Characteristics:* | | | | |
| $\overline{\Delta}_i$ | 0.383 | -0.020 | -0.006 | 0.067 |
| $\sigma_{\Lambda_i}$ | 0.517 | 0.676 | 0.271 | 0.407 |
| $\overline{d}_i$ | 87.925 | 54.723 | 241.400 | 170.601 |
| *Sectoral Compositions:* | | | | |
| % in Retail | 0.389 | 0.379 | 0.319 | 0.346 |
| % in Service | 0.416 | 0.424 | 0.595 | 0.523 |
| *Location Compositions:* | | | | |
| % in Manhattan | 0.457 | 0.567 | 0.520 | 0.518 |
| % in Brooklyn | 0.438 | 0.340 | 0.335 | 0.356 |
| *Regression Estimates:* | | | | |
| $\rho$ | 0.422*** | 0.164 | 0.775*** | 0.727*** |
| | [0.145] | [0.152] | [0.073] | [0.082] |
| $\sigma_{\epsilon_{i,t}}$ | 0.596 | 0.435 | 0.276 | 0.411 |
| *Arellano-Bond Test for:* | | | | |
| $AR(4)P - value$ | 0.963 | 0.349 | 0.019 | 0.316 |
| $AR(5)P - value$ | 0.210 | 0.361 | 0.171 | 0.905 |
| *Fraction of Establishments* | 0.194 | 0.220 | 0.586 | 1.000 |

*Note:* Standard errors are reported in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Robustness Check: Characteristics & Estimates of Clusters

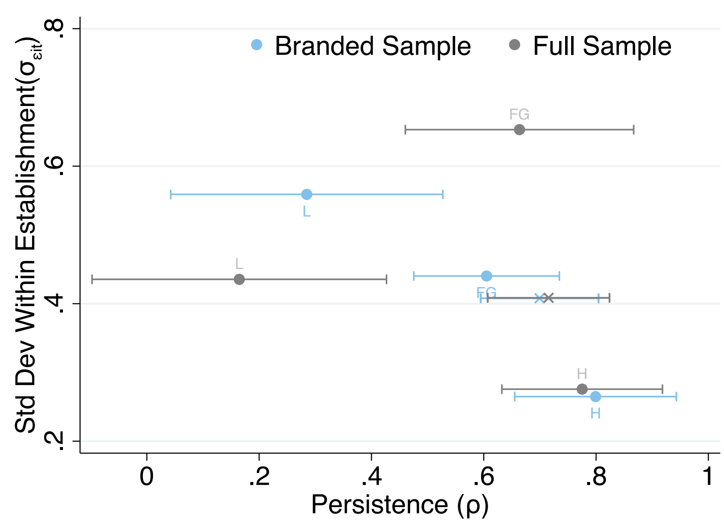


*Note:* The figure plots persistence estimates and within-establishment standard deviations for the brand sample (blue) and full sample (gray). "x" marks indicate pooled regression estimates for each sample. Horizontal lines denote 95% confidence intervals for persistence. Cluster labels are abbreviated as "L" for Low-Traffic, "FG" for Fast-Growing, and "H" for High-Traffic.

Figure 8: Robustness Check: Clusters and Pooled Estimates

Then, we estimate the process as outlined in equations (3) and (5) for the year-over-year demand growth rate of all establishments, including those without brands, in each cluster. Observations in the full sample are weighted by the ratio of the

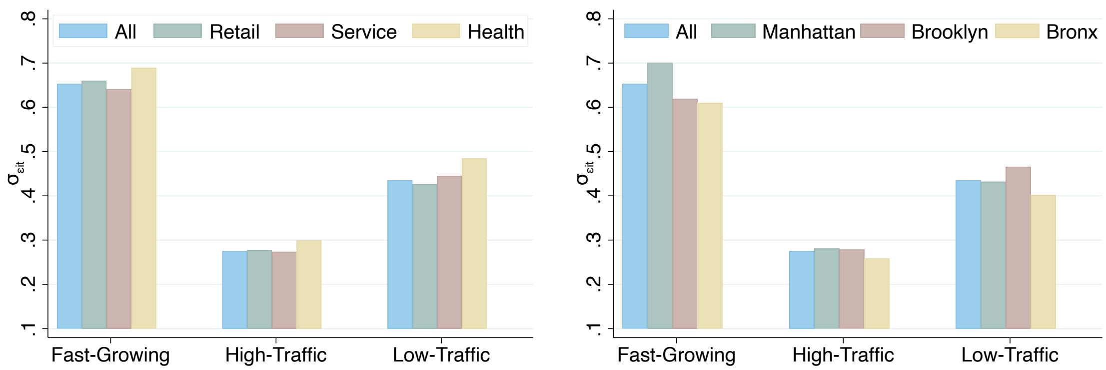


(a) By Sectors

(b) By Boroughs

*Note:* The figure shows the standard deviation of residuals for establishments in each cluster from the full sample, categorized by sectors or boroughs. The horizontal axis in both panels lists cluster groups. Blue bars show the overall cluster-level standard deviation (see Table 3); other bars show the same statistic for each sector or borough within a cluster.

Figure 9: Robustness Check: Standard Deviation of Residuals By Groups

borough-sector-specific number of establishments in the brand sample to those in the full sample. Table 3 summarizes the regression estimates. Figure 8 presents these estimates with 95% confidence intervals, comparing them to those from the brand sample, which confirm the robustness of our persistence estimates. Note that in the full sample, the Fast-Growing group exhibits the highest standard deviation of residuals, while in the brand sample, the Low-Traffic group shows the highest standard deviation.

Similar to Figure 7, Figure 9 also displays the standard deviation of the residuals for establishments within each cluster by sectors or boroughs, using data from the full sample. Based on Figure 9, there is no evidence that location or sectoral differences explain the heterogeneity in demand shocks across clusters.

# 6 Why Does This Matter?

A large literature in macroeconomics and industrial organization models firm behavior under the assumption that all establishments draw productivity or demand from a *single* stochastic process. In canonical firm dynamics models, such as Hopenhayn (1992) and Hopenhayn and Rogerson (1993), firms evolve according to a common AR(1) process, single cutoff that governs decisions like entry, exit, hiring, or investment. These models typically rely on a single productivity distribution, so policy changes affect firms uniformly through shifts in this threshold. Even richer environments incorporating financial frictions (e.g., Midrigan and Xu (2014)) or imperfect competition (e.g., Asplund and Nocke (2006)) maintain this single-process structure. Many models of firm dynamics continue to build on this foundation, linking discrete choices—such as opening new establishments (e.g. Cao et al. (2017) or Smith and Ocampo (2025)) or hiring and firing (e.g. Alvarez and Veracierto (2001))— to the shared underlying process. This simplifying assumption also appears in models of technology adoption (e.g., Mukoyama (2003); Buera and Shin (2013))), where firms face fixed costs to upgrade technology but still draw productivity from a common stochastic process. While recent research highlights the importance of realistic shock modeling—e.g., Jaimovich et al. (2025) show that deviations from Gaussian AR(1) shocks can materially alter macro outcomes—these models typically retain the assumption of a single underlying shock process.

In contrast, our evidence demonstrates that establishments face multiple, systematically different demand processes. This richer structure implies that models relying on a single cutoff may misidentify marginal firms and mischaracterize how policy propagates through the cross-section of firms. By quantifying this cluster-level heterogeneity, our framework offers a data-driven foundation for calibrating dynamic heterogeneous-firm models—including those used for policy analysis and business-cycle research—much in the spirit of structural calibration approaches like Fernández-Villaverde et al. (2016).

## 6.1 A Simple Dynamic Model with Cutoff Decisions

We present a highly stylized dynamic model in which an establishment chooses whether to incur a cost (e.g., to invest, expand capacity, hire additional workers, or open a new establishment) based on its demand-growth state. The goal is not to match the richness of firm dynamics, but to illustrate how different stochastic processes for demand shocks imply different marginal establishments, even when all other primitives are held constant. To keep the problem analytically tractable, we treat the upgrade decision as a once-and-for-all choice in the current period, abstracting from the option value of delaying the investment.[10]

Let the state variable be the establishment's year-over-year demand growth,

$$s_t = \tilde{\Delta}_{i,t},$$

which evolves according to an AR(1) process similar to that in eq. (5). Our objective is to contrast the standard 'single-process benchmark' typically used in the literature with a 'three-process specification' motivated by our empirical results.

**Single-Process Benchmark.** In the canonical heterogeneous-firm environment, all establishments draw shocks from the same AR(1):

$$s_{t+1} = \rho s_t + \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim N(0, \sigma^2).$$

Suppose that per-period profits depend linearly on demand growth,

$$\pi^0(s_t) = \alpha + \gamma s_t,$$

[10] In canonical heterogeneous-firm models, establishments re-optimize in every period and may delay upgrading until the continuation value exceeds the associated fixed cost; see Hopenhayn (1992), Hopenhayn and Rogerson (1993), and Midrigan and Xu (2014). This formulation introduces a nonlinearity through the $\max\{V_0, V_1\}$ operator in the Bellman problem, eliminating closed-form solutions. Our once-and-for-all assumption abstracts from this option value and allows the cutoff to be derived analytically while preserving the core comparative-static insight.

and that upgrading increases the slope (permanently) to $\gamma_1 = (1+\phi)\gamma$ with $\phi > 0$:

$$\pi^1(s_t) = \alpha + \gamma_1 s_t.$$

The firm decides whether to upgrade or not in period 0. If it chooses to upgrade, it commits to pay $F$ every period from then on in exchange for higher production capacity (e.g. access to the $\pi_1$ technology). The value of not upgrading is denoted by $V_0(s)$ and the value after upgrading by $V_1(s)$. The Bellman equations are

$$V_0(s) = \pi_0(s) + \beta\mathbb{E}[V_0(s') \mid s] \qquad \text{and} \qquad V_1(s) = \pi_1(s) - F + \beta\mathbb{E}[V_1(s') \mid s].$$

Because the profit function and the transition law are linear, the value functions are themselves linear, satisfying:

$$V_0(s) = \frac{\alpha}{1-\beta} + \frac{\gamma}{1-\beta\rho}\, s, \qquad V_1(s) = \frac{\alpha - F}{1-\beta} + \frac{\gamma_1}{1-\beta\rho}\, s.$$

The upgrade decision $V_1(s) \geq V_0(s)$ yields the unique cutoff

$$s^\star = \frac{F(1-\beta\rho)}{(1-\beta)(\gamma_1 - \gamma)} = \frac{F(1-\beta\rho)}{(1-\beta)\phi\gamma}.$$

Under a common shock process, all establishments therefore share the same cutoff $s^\star$. This is the standard mechanism in heterogeneous-firm models: a single set of marginal establishments governs entry, exit, exporting, or investment responses to policy.

**Three Distinct Shock Processes.** Motivated by our empirical findings, suppose instead that establishments belong to one of three groups

$$g \in \{H, G, L\},$$

each characterized by a distinct AR(1) demand-growth process,

$$s_{t+1} = \rho_g s_t + \epsilon^g_{t+1}, \qquad \epsilon^g_{t+1} \sim N(0, \sigma^2_g).$$

We can solve for the Bellman equations following similar steps to those in the benchmark case, and replacing $\rho$ with $\rho_g$. The closed-form value functions are

$$V^g_0(s) = \frac{\alpha}{1-\beta} + \frac{\gamma}{1-\beta\rho_g}\, s, \qquad V^g_1(s) = \frac{\alpha - F}{1-\beta} + \frac{\gamma_1}{1-\beta\rho_g}\, s.$$

The cutoff determining group-$g$'s cutoff is therefore

$$s^\star_g = \frac{F(1-\beta\rho_g)}{(1-\beta)\phi\gamma},$$

which is strictly decreasing in $\rho_g$. Because $\rho_H > \rho_G > \rho_L$ in the data, the cutoffs satisfy

$$s^\star_L > s^\star_G > s^\star_H.$$

Establishments in the High-Traffic group (high persistence) upgrade at the lowest $s$, while those in the Low-Traffic group require substantially stronger current demand growth.

**Why the Single-Process Benchmark Is Misleading.** The single-process benchmark implies a *single* marginal establishment, with policy responses governed by the cutoff $s^\star$. In reality, the continuation value $\mathbb{E}[V_g(s') \mid s]$ is increasing in persistence $\rho_g$, so ignoring heterogeneity collapses three distinct continuation-value mappings into one. This leads to the following two failures:

1. *Misidentification of marginal establishments*. Using a common, single shock process misidentifies the marginal establishments that are indifferent between upgrading and not upgrading. Because $\rho_H > \rho > \rho_G > \rho_L$ in the data, the single process cutoff $s^\star$ lies between the true group specific cutoffs,

$$s_L^\star > s_G^\star > s^\star > s_H^\star.$$

As a result, the model incorrectly predicts that Low Traffic and Fast Growing establishments upgrade when they should not, while failing to capture upgrades by High Traffic establishments that are in fact marginal.

2. *Incorrect comparative statics*. The single shock benchmark also generates incorrect comparative statics. When policies change $F$, $\phi$, or $\beta$, the cutoff $s_g^\star$ responds differently across groups because the change in the upgrade threshold depends on the persistence $\rho_g$. Establishments facing less persistent demand shocks, such as those in the Low Traffic and Fast Growing groups, exhibit larger changes in their upgrade cutoffs, while establishments subject to more persistent shocks, namely those in the High Traffic group, respond more weakly. Imposing a common shock process suppresses this heterogeneity and forces all establishments to respond as if they shared the same degree of persistence.

Aggregation therefore produces biased predictions for outcomes such as investment, exporting, capacity expansion, and job creation, depending on the specific application. In short, even in this extremely simple environment, heterogeneous shock processes imply heterogeneous marginal establishments. Imposing a single-process structure, as is standard in heterogeneous-firm models, therefore obscures the true propagation of shocks and policy across consumer-facing industries.

## 6.2 Numerical Example

To illustrate the implications of our findings, consider a policy that loosens regulations and reduces the cost of adding a new establishment—such as a regulatory reform that reduces entry barriers $F$. We use this example to compare outcomes under two different assumptions about the demand process: one in which all establishments face a common stochastic process for demand growth, and another in which establishments face three systematically distinct demand processes, as supported by our empirical evidence. This quantitative exercise highlights how model predictions—particularly for investment responses—can differ substantially depending on whether demand heterogeneity is accounted for.

To that end, we simulate 30,000 firms, with group shares reported in Table 2 (contained in the row 'fraction of establishments'). The baseline parameters are $\beta = 0.98$, $\phi = 0.50$, $\alpha = 1.00$, and $\gamma = 1.00$, while the fixed cost $F$ moves from 0.01 to 0.02. Figure 10 reports the share of firms that upgrade (that is, add an establishment) and the average investment per firm. Aggregate investment is this value multiplied by the number of firms in the economy. The figure presents results for both the single-shock process and the three distinct shock processes.

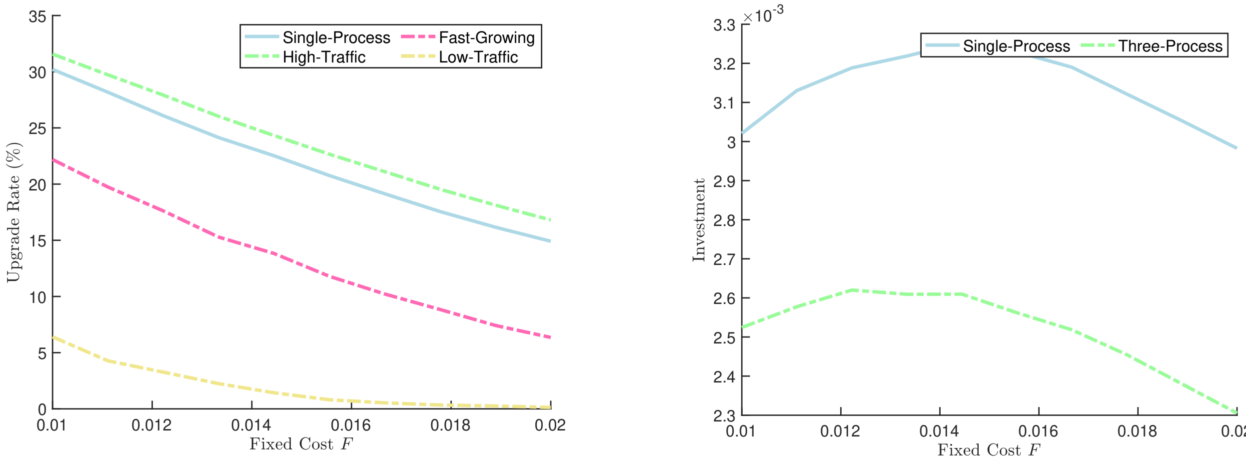


(a) Upgrade Percentage

(b) Investment

Figure 10: Change in Fixed Cost $F$

With only a single, common stochastic process, one would predict more investment and more firms upgrading than actually occurs at every level of $F$, as shown in Figure 10. Although a larger share of establishments in the high traffic group choose to upgrade and invest more than predicted by the single process, they account for only 58.8 percent of all establishments. In contrast, a smaller share of establishments in the Fast Growing and Low Traffic groups upgrade, and their investment levels fall below those implied by the single process. As a result, the single process benchmark generates aggregate investment that exceeds the level implied by the three distinct shock processes. This could, of course be abrogated by changing the calibration, but would then lead one to incorrectly identify other parameters. For instance, if $\phi = 1$ is the true gain in productivity from an investment, with the single, common productivity process, we would infer that $\phi$ is about $\frac{1}{4}$ lower.

A similar experiment can be conducted to measure the effects of innovation subsidies, which we represent by increasing $\phi$ or monetary policy contraction, increasing the discount rate by decreasing $\beta$. Figure 11 shows how these policies affect upgrading decisions under the single shock and under the heterogeneous shock process environment. In both cases, policy changes generate heterogeneous responses across establishment groups that are obscured by the single process specification.

An increase in the innovation subsidy raises upgrade rates across all groups by lowering the relevant upgrade cutoffs.[11] When these heterogeneous responses are aggregated, the single-shock benchmark again predicts higher aggregate investment than the heterogeneous-shock process environment. A similar pattern emerges under changes in monetary policy: a lower discount factor increases upgrade incentives in all groups, but the magnitude of the response depends on the shock persistence, and

[11] Although policy changes shift cutoffs by different amounts across groups, larger cutoff movements do not necessarily translate into larger changes in upgrade shares. what matters is the density of establishments around the cutoff. Because the state variable follows a normal distribution, groups with more mass near the margin may exhibit larger responses even when their cutoff shifts are smaller.

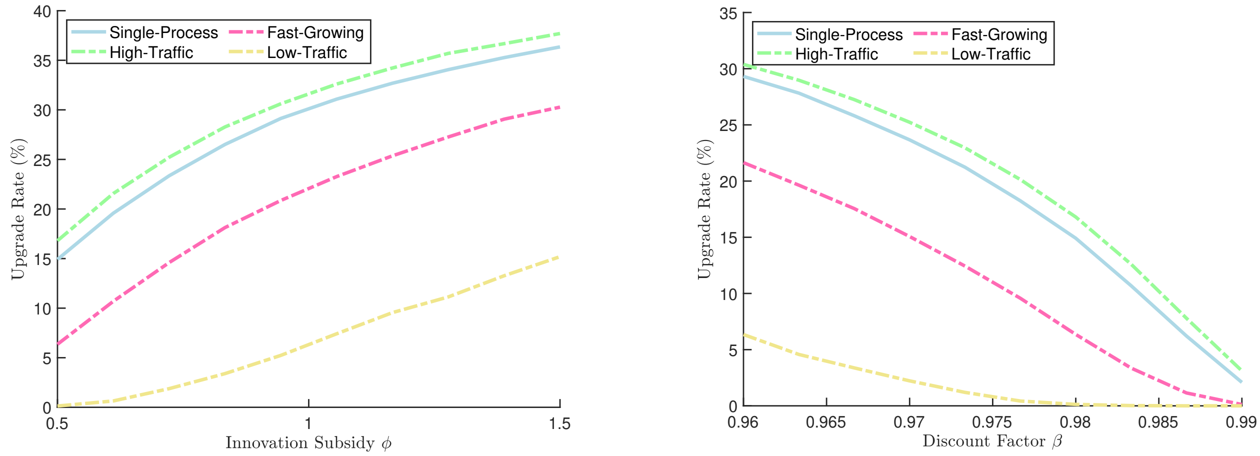


(a) Fraction upgrading by $\phi$

(b) Fraction upgrading by $\beta$

Figure 11: Changes in the investment threshold and policy

the single-shock benchmark systematically overstates the resulting increase in aggregate investment.

# 7 Conclusion

We analyze demand dynamics among various groups of branded establishments in customer-facing retail, services and health sectors using SafeGraph's monthly foot-traffic data for the pre-pandemic period from 2018-2019. Our high-frequency panel data allow us to explore both the annual growth and monthly volatility in demand.

To estimate demand shocks using establishment foot-traffic data, we begin by eliminating location-specific biases and trend biases in sampling. We obtain the adjusted number of visitors for each establishment every month. Then, we find establishments with high foot traffic still exhibit considerable within-year variation. Besides, there is heterogeneity in the foot traffic growth pattern. Although monthly volatility occurs in establishments of all sizes, it is substantial in establishments in the two tails of the average annual growth rate distribution in foot traffic. Thus, growing or declining paths are not smooth over months.

We develop a theoretical framework to explore the relationship between foot traffic and actual demand. Changes in foot traffic may also reflect firms' endogenous decisions, such as promotional strategies or advertising campaigns. To address this, we construct a model that isolates establishment-level demand shocks from firms' endogenous actions. This theoretical model provides the foundations for using changes in foot traffic to estimate demand shocks in the empirical analysis.

We apply a $k-$means clustering method to discuss the demand growth heterogeneity more systematically. In this paper, we categorize establishment clusters based on establishment-specific average foot traffic, the year-over-year growth rate, and monthly volatility in demand. We obtain three distinct groups: Fast-Growing, Low-Traffic, and High-Traffic clusters. The Fast-Growing cluster, comprising 25.6% of establishments, exhibits the highest annual growth in foot traffic. Most establishments in this group operate in the retail and service sectors and are primarily located in Manhattan and Brooklyn. The Low-Traffic cluster, representing 15.5% of establishments, is characterized by the lowest average monthly foot traffic. Within this group, 53.4% belong to the retail sector, with 68.1% located in Manhattan. The High-Traffic cluster contains establishments with the highest average monthly foot traffic, with 59% operating in the service sector, and most are also based in Manhattan.

We then estimate the AR(1) process for the year-over-year foot-traffic growth rate with MA(4) errors for all branded establishments in different clusters. By comparing the estimates for each cluster and the pooled regression results, we conclude that the analysis of the average effect is not representative of the actual data observed in business due to considerable heterogeneity in estimates across establishments clusters. Establishments in the High-Traffic cluster consistently experience the lowest variability in demand shocks across all boroughs and sectors.

Our analysis reveals that demand dynamics in customer-facing industries are fundamentally heterogeneous, and this heterogeneity is driven largely by establishment-specific factors—differences in persistence, innovation variance, and underlying growth trends—rather than by sector, geography, or brand-level strategies. This insight overturns the common assumption, implicit in much empirical and theoretical work, that demand shocks can be treated as homogeneous or pooled across firms within a sector or region. Instead, establishments exhibit distinct demand processes, reflecting micro-level fundamentals such as managerial practices, local competition, neighborhood characteristics, and other unobserved attributes. Importantly, this variation is not entirely idiosyncratic: our clustering analysis identifies three stable and economically meaningful groups—Fast-Growing, Low-Traffic, and High-Traffic—each with its own characteristic shock process. Recognizing this structure is essential. Models that ignore these clusters and rely on pooled or sector-level shocks risk conflating fundamentally different demand environments and misrepresenting both the persistence and volatility of shocks. By identifying and quantifying this cluster-level heterogeneity, our framework provides a more accurate and tractable foundation for modeling demand in customer-facing sectors and improves the external validity of applications that rely on establishment-level demand shocks.

# References

**Alvarez, Fernando and Marcelo Veracierto**, "Severance payments in an economy with frictions," *Journal of Monetary Economics*, 2001, *47* (3), 477–498.

**Asplund, Marcus and Volker Nocke**, "Firm turnover in imperfectly competitive markets," *Review of Economic Studies*, 2006, *73* (2), 295–327.

**Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa**, "Discretizing unobserved heterogeneity," *Econometrica*, 2022, *90* (2), 625–643.

**Buera, Francisco J. and Yongseok Shin**, "Financial Frictions and the Persistence of History: A Quantitative Exploration," *Econometrica*, 2013, *81* (3), 888–931.

**Cao, Dan, Henry R. Hyatt, Toshihiko Mukoyama, and Erick Sager**, "Firm Growth Through New Establishments," Working Paper 3361451, SSRN February 2017.

**Cronin, Christopher J and William N Evans**, "Private Precaution and Public Restrictions: What Drives Social Distancing and Industry Foot Traffic in the COVID-19 Era?," Working Paper 27531, National Bureau of Economic Research July 2020.

**Farboodi, Maryam, Gregor Jarosch, and Robert Shimer**, "Internal and external effects of social distancing in a pandemic," *Journal of Economic Theory*, 2021, *196*, 105293.

**Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, and Frank Schorfheide**, "Solution and Estimation Methods for DSGE Models," *Handbook of Macroeconomics*, 2016, *2*, 527–724.

**Goldfarb, Avi and Catherine Tucker**, "Which Retail Outlets Generate the Most Physical Interactions?," Working Paper 27042, National Bureau of Economic Research April 2020.

**Goolsbee, Austan and Chad Syverson**, "Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020," *Journal of Public Economics*, 2021, *193*, 104311.

**Gourio, Francois and Leena Rudanko**, "Customer capital," *Review of Economic Studies*, 2014, *81* (3), 1102–1136.

**Haltiwanger, John, Ron S. Jarmin, Robert Kulick, and Javier Miranda**, "High Growth Young Firms: Contribution to Job, Output, and Productivity Growth," in "Measuring Entrepreneurial Businesses: Current Knowledge and Challenges" NBER Chapters, National Bureau of Economic Research, Inc, december 2016, pp. 11–62.

**Hopenhayn, Hugo A.**, "Entry, exit, and firm dynamics in long run equilibrium," *Econometrica*, 1992, *60* (5), 1127–1150.

**_ and Richard Rogerson**, "Job turnover and policy evaluation: A general equilibrium analysis," *Journal of Political Economy*, 1993, *101* (5), 915–938.

**Inoue, Atsushi and Gary Solon**, "A portmanteau test for serially correlated errors in fixed effects models," *Econometric Theory*, 2006, *22* (5), 835–851.

**Jaimovich, Nir et al.**, "The Empirical Distribution of Firm Dynamics and its Macro Implications," https://www.nirjaimovich.com/s/Shocks_Matter_January2025.pdf 2025. Working paper.

**__ , Stephen J Terry, and Nicolas Vincent**, "The Empirical Distribution of Firm Dynamics and Its Macro Implications," Working Paper 31337, National Bureau of Economic Research June 2023.

**Kaplan, Greg and Guido Menzio**, "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," *Journal of Political Economy*, 2016, *124* (3), 771–825.

**Liang, Yunlei, Song Gao, Yuxin Cai, Natasha Zhang Foutz, and Lei Wu**, "Calibrating the dynamic Huff model for business analysis using location big data," *Transactions in GIS*, 2020, *24* (3), 681–703.

**Melcangi, Davide and Silvia Sarpietro**, "Nonlinear Firm Dynamics," *FRB of New York Staff Report*, 2024, (1088).

**Midrigan, Virgiliu and Daniel Yi Xu**, "Finance and misallocation: Evidence from plant-level data," *Econometrica*, 2014, *82* (2), 529–576.

**Mukoyama, Toshihiko**, "Innovation, Imitation, and Growth with Cumulative Technology," *Journal of Monetary Economics*, 2003, *50* (2), 361–380.

**Petrosky-Nadeau, Nicolas and Etienne Wasmer**, "Macroeconomic dynamics in a model of goods, labor, and credit market frictions," *Journal of Monetary Economics*, 2015, *72* (C), 97–113.

**Smith, Dominic A. and Sergio Ocampo**, "The Evolution of US Retail Concentration," *American Economic Journal: Macroeconomics*, January 2025, *17* (1), 71–101.

**Sterk, Vincent, Petr Sedláček, and Benjamin Pugsley**, "The nature of firm growth," *American Economic Review*, 2021, *111* (2), 547–79.

**Storesletten, Kjetil, Jose-Victor Rios Rull, and Yan Bai**, "Demand Shocks that Look Like Productivity Shocks," 2011 Meeting Papers 99, Society for Economic Dynamics 2011.

**Syverson, Chad**, "What Determines Productivity?," *Journal of Economic Literature*, June 2011, *49* (2), 326–65.

# A Adjusting "Raw" Foot Traffic

This section documents the sampling biases in the raw foot-traffic data and details the adjustment procedures implemented to correct them.

## A.1 Sampling Bias in "Raw" Foot Traffic

Figure A12 shows monthly raw foot traffic overall and by borough. In panel (a), total foot traffic more than doubles in the first four months of observation before stabilizing. Panel (b) highlights heterogeneity across boroughs: the Bronx experiences a markedly larger early increase compared with Brooklyn and Manhattan.

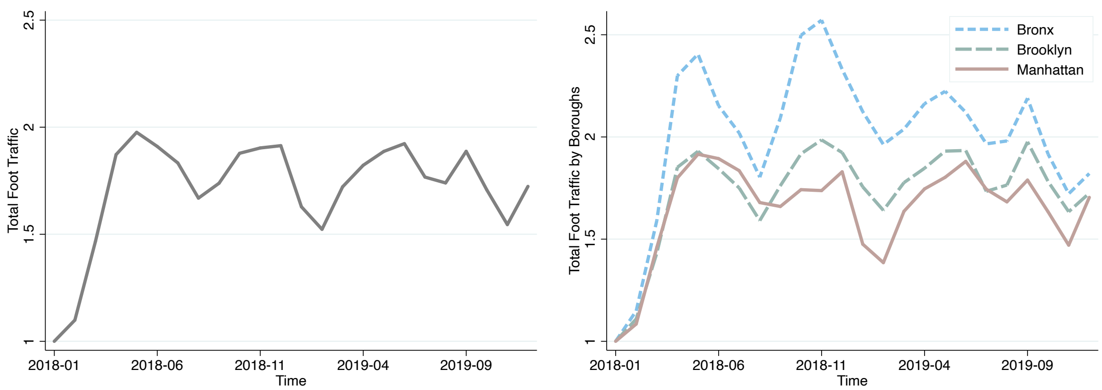


(a) All Establishments (b) By Borough

*Note:* Monthly foot traffic for each POI is computed by summing visitor counts from all home CBGs, then aggregating across POIs. The first observation is normalized to 1 to highlight the magnitude of the initial increase, both overall and by borough.

Figure A12: Normalized Raw Foot Traffic Over Time

## A.2 Adjustment Process

We denote each establishment (POI) as $i$. The total number of unique POIs within borough $c$ at month $t$ is denoted as $I_{c,t}$. We only consider three boroughs in this paper, so $c \in \{\text{NYC, Brooklyn, Bronx}\}$. From the *Place Patterns* data, we observe the home residence at the CBG level of consumers who visit POI $i$ located in borough $c$ during month $t$. Their home CBGs are denoted by $\ell \in \{1, ..., \mathcal{L}_{c,i,t}\}$, where $\mathcal{L}_{c,i,t}$ presents the total number of unique home CBGs of customers visiting POI $i$ located in borough $c$ during month $t$. The total number of unique visitors who reside in $\ell$ and visit establishment $i$ located in borough $c$ at month $t$ is denoted by $V_{c,i,\ell,t}$. Hence, foot traffic in all POIs located in borough $c$ at month $t$ is denoted by $\bar{V}_{c,t}$, and

$$\bar{V}_{c,t} = \sum_{i=1}^{I_{c,t}} \sum_{\ell=1}^{\mathcal{L}_{c,i,t}} V_{c,i,\ell,t}.$$

Total foot traffic in month $t$ is denoted by $FT_t$, and

$$FT_t = \sum_{c} \bar{V}_{c,t}.$$

In Figure A12, the left panel shows the total foot traffic $FT_t$ over time, and the right panel depicts total foot traffic for three boroughs $\bar{V}_{c,t}$ over time.

**Addressing Location-specific Biases** We collect the total number of devices located in all home CBGs during a specific time from the *Home Panel Summary*.[12] We denote the total number of devices located in CBG $\ell$ at time $t$ by $S_{\ell,t}$. Subsequently, we calculate the proportion of devices in our sample relative to the local total population in 2019. We obtain the the total population data from the American Community Survey (ACS) and denote it by$P_\ell$ for CBG $\ell$. We compute the location-bias adjustment factors such that $\bar{V}_{c,t}$, the total foot traffic to establishments in each borough at $t$, remains unchanged,

$$\bar{V}_{c,t} = \sum_{i=1}^{I_{c,t}} \sum_{\ell=1}^{\mathcal{L}_{c,i,t}} V_{c,i,\ell,t} \frac{TS_{c,t}}{S_{\ell,t}/P_\ell}. \tag{6}$$

We denote the location-bias adjustment factors each CBG $\ell$ in borough $c$ at month $t$ as $A_{c,\ell,t}$, which satisfies

$$A_{c,\ell,t} = \frac{TS_{c,t} * P_\ell}{S_{\ell,t}},$$

with $TS_{c,t}$ defined such that equation (6) holds with equality,

$$TS_{c,t} = \frac{\sum_{i=1}^{I_{c,t}} \sum_{\ell=1}^{\mathcal{L}_{c,i,t}} V_{c,i,\ell,t}}{\sum_{i=1}^{I_{c,t}} \sum_{\ell=1}^{\mathcal{L}_{c,i,t}} \frac{V_{c,i,\ell,t} * P_\ell}{S_{\ell,t}}}.$$

While this adjustment fixes cross-sectional biases, it still retains the aggregate trend described in Figure A12.

**Addressing Trend Biases** To correct for trends, we divide the number of unique visitors adjusted for location-specific bias, $V_{c,i,\ell,t}A_{c,\ell,t}$, by the total foot traffic in our sample, $FT_t$, and normalize it by the long-run average of foot traffic $\overline{FT}$ ($\overline{FT} = \frac{\sum_t FT_t}{24}$). Hence, the number of unique visitors adjusted for both location-specific bias and time trends is denoted by

$$\hat{V}_{c,i,\ell,t} = V_{c,i,\ell,t} A_{c,\ell,t} \frac{\overline{FT}}{FT_t}.$$

Notice that we compute the location-bias adjustment factors such that the total raw foot traffic to POIs in each borough at time $t$ remains unchanged, so

$$\sum_{i=1}^{I_{c,t}} \sum_{\ell=1}^{\mathcal{L}_{c,i,t}} \hat{V}_{c,i,\ell,t} = \frac{\bar{V}_{c,t}}{FT_t} \overline{FT}.$$

Figure A13 displays the evolution of adjusted foot traffic over time for each borough as a proportion of $\overline{FT}$. Each line represents the monthly percentage of foot traffic within a borough relative to the total foot traffic across all boroughs. The summation of these proportions for all boroughs in a given month equals one. This adjustment notably eliminates the trends observed in Figure A12. In what follows, we use adjusted foot traffic $\hat{V}_{c,i,\ell,t}$ for the analysis.

[12]Specifically, we use the number of devices residing in each CBG from the *Home Panel Summary* data. This variable captures the number of distinct devices observed with a primary nighttime location in the specified census block group.

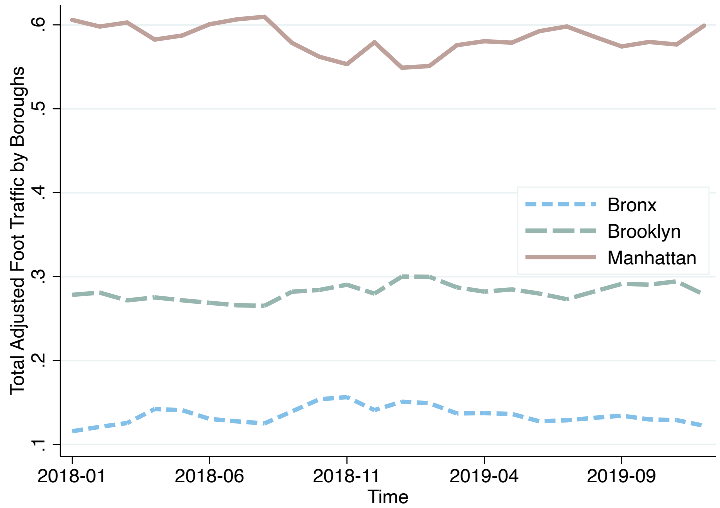


Figure A13: Adjusted Foot Traffic

**Adjusted Foot Traffic** Figure A13 plots adjusted foot traffic by borough, expressed as each borough's share of total foot traffic across all boroughs. By construction, the monthly shares sum to one. The adjustments remove the artificial trends evident in Figure A12 and reveal the relative scale of foot traffic across boroughs.

# B Cleaning Process

In this section, we show the steps used in the cleaning process and their effects on the sample size. Table B4 shows how each step reduces the sample size. From our initial set, the first large reduction is to only take large enough POIs. That is, in each period, if a POI has fewer than 10 visitors in the month, we drop it. Note that doing this first implies that it will also restrict our panel, because our restriction on length then becomes having a sufficient number of months with a sufficient number of visitors. The rest of our selection mostly takes out extreme fluctuations that are due to measurement error until our final restriction to include only the POIs with an associated brand.

Figure B14 illustrates the ratio of total monthly adjusted foot traffic between the brand sample and the full sample. We use the brand sample for further analyses in this paper.

# C The Branded Sample

## C.1 Descriptive Statistics of Establishment and Foot Traffic

Figure C20 provides a more detailed breakdown of the top 10 (out of 34) SafeGraph-defined categories, which together represent about 88% of unique establishments in

| | # of POIs | % of POIs | # of Obs. |
|---|---|---|---|
| Original POIs all boroughs | 71,683 | 100.00% | 1,513,157 |
| Drop if < 10 monthly visitors | 67,809 | 94.60% | 1,348,087 |
| Drop if YoY change does not exist | 60,811 | 84.83% | 1,262,762 |
| Keep Retail, Services, and Health | 47,678 | 66.51% | 1,020,406 |
| Drop POIs w/ extreme fluctuations | 47,248 | 65.91% | 1,011,498 |
| Drop POIs w/ extreme YoY growth rate | 47,201 | 65.85% | 1,010,620 |
| Drop POIs with top 1% average YoY growth rate over time | 46,729 | 65.19% | 1,002,706 |
| Drop POIs w/ < two MoM change in visitors | 44,011 | 61.40% | 993,164 |
| | | | *(Full Sample)* |
| Keep branded POIs | 5,924 | 8.26% | 137,512 |
| | | | *(Branded Sample)* |

Table B4: Sample Selection

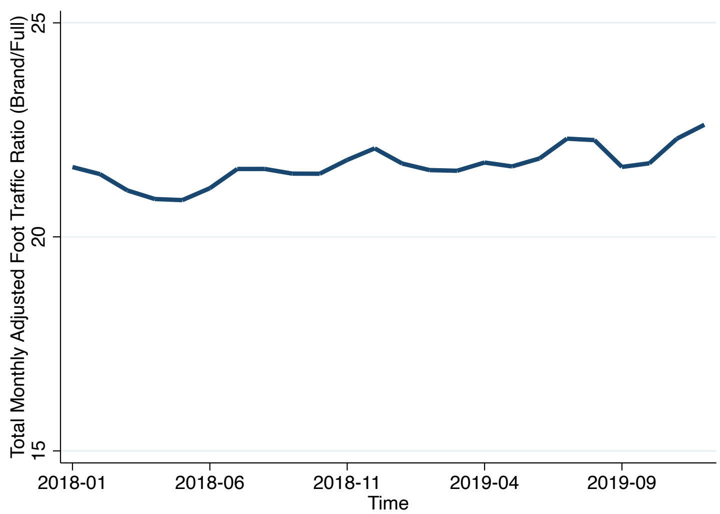


Figure B14: Total Adjusted Foot-Traffic Ratio

the branded sample.[13] The largest single group is restaurants and eating places, which align closely with the demand framework we study: food production is recorded as revenue only when customers physically visit the establishment.

Figure C16 shows that most establishments are corporate-owned or franchises with limited autonomy over pricing and advertising decisions that could endogenously affect foot traffic. This can be validated anecdotally looking at Table C1, which

[13]SafeGraph-defined categories are more granular than NAICS codes for understanding specific place types.

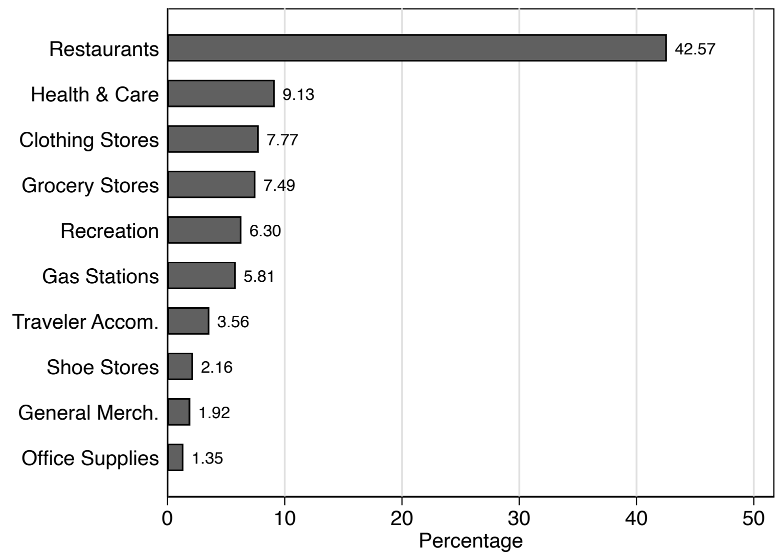


Figure C15: Top 10 Categories of Establishments

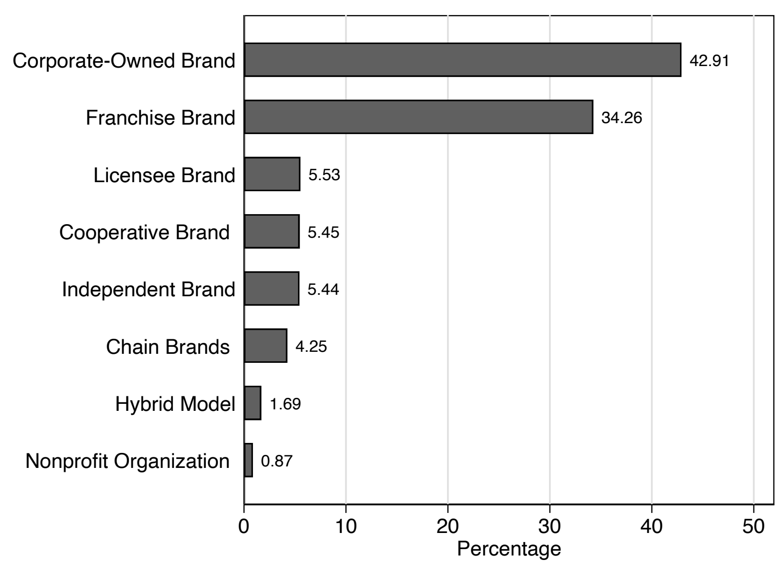


Figure C16: Establishments by Brand Types

| Brand Name | Brand Type | Share (%) |
|---|---|---|
| Dunkin' | Franchise Brand | 6.20 |
| Starbucks | Corporate-Owned Brand | 4.19 |
| Sweetgreen | Corporate-Owned Brand | 3.71 |
| Subway | Franchise Brand | 2.62 |
| Walgreens | Corporate-Owned Brand | 2.55 |
| Baskin Robbins | Franchise Brand | 2.30 |
| McDonald's | Franchise Brand | 2.26 |
| BP | Licensee Brand | 1.92 |
| Duane Reade by Walgreens | Corporate-Owned Brand | 1.42 |
| CVS | Corporate-Owned Brand | 1.40 |

*Note:* This table presents the top 10 brands of establishments in the branded sample, their type, and their unweighted share of all branded establishments.

Table C1: Top 10 Brands of Establishments

shows the top 10 brands among the establishments. This can be validated anecdotally looking at Table C1, which shows the top 10 brands among the establishments. Most top brands in the sample are corporate or franchise chains with centralized marketing, standardized offerings, and largely brand-wide prices.

Table C2 outlines the distribution of establishments by sector and location within our sample. Each cell in the table provides two ratios: the first indicates the proportion of establishments within a sector for each location, and the second shows the distribution of sectors within a borough. The table reveals that the majority of Manhattan establishments are in the service sector, whereas in Brooklyn and the Bronx, retail dominates. Additionally, Manhattan hosts the majority of establishments across all sectors compared to the other boroughs.

| | Retail | Service | Health |
|---|---|---|---|
| Manhattan | 47%/37% | 63%/60% | 43%/3% |
| Brooklyn | 33%/51% | 25%/46% | 28%/3% |
| Bronx | 20%/55% | 12%/39% | 29%/6% |

Table C2: Distribution of POIs Across Locations and Sectors (%)

Figure C17 presents the average number of foot traffic among POIs in each sector for a given month both across years (left panel) and by different years (right panel). The average foot-traffic trends for each sector are similar in 2018 and 2019. On average, establishments in the service sector have higher foot traffic than other sectors for all months. The average foot traffic for POIs in all sectors slightly decline in the first half of the year but then rebound to the original level in the latter half of the year. So, we conclude the average foot-traffic processes for all sectors are stationary cross time.

Figure C18 displays the distribution of POIs on average foot traffic across time. We exclude those POIs with top 1% average foot traffic, but the distribution is still right-skewed.[14]

[14] Those POIs with the top 1% average foot traffic are mainly in the retail (53%) and service sectors (47%) and in Manhattan (66%). Particularly, *Restaurants and Other Eating Places*, *Department Stores*, and *Health and Personal Care Stores* are the three top categories.

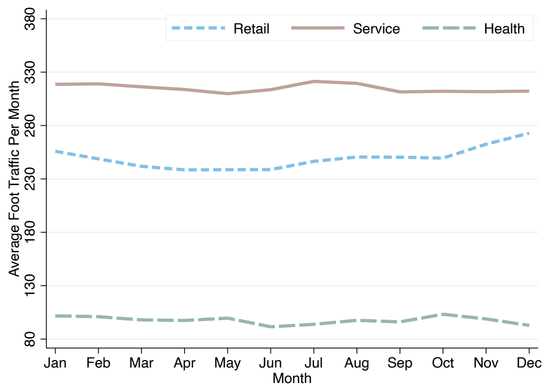


(a) Across Years

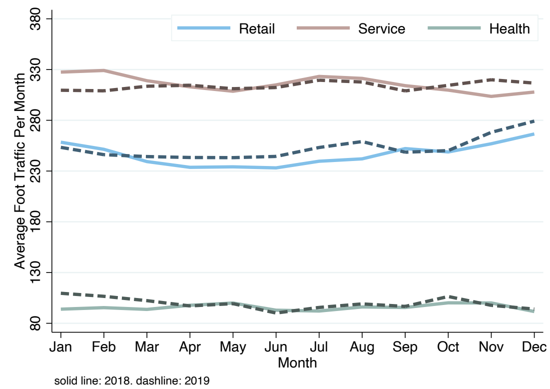


(b) By Years

Figure C17: Average Number of Visitors per POI

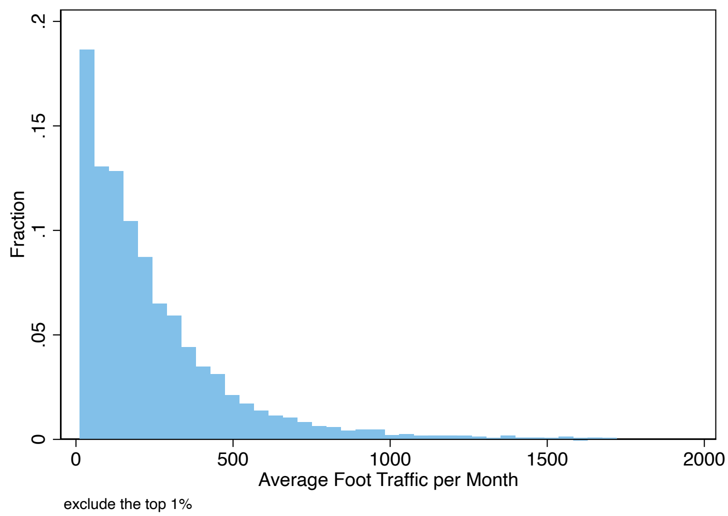


Figure C18: Distribution of Establishment on Average Monthly Foot Traffic

## C.2 Foot Traffic Growth Rates

We categorized POIs by their average monthly foot traffic into ten deciles. Figure C19 illustrates the average growth rate and range of monthly volatility in foot traffic for POIs in each decile. In Figure C19, the left panel demonstrates that establishments with higher monthly foot traffic exhibit lower average year-over-year growth rates in foot traffic, suggesting the fast-growing establishments with fewer customers every month. The right panel illustrates increasing average foot traffic by deciles (blue dots) and shows the 90th, 50th, and 10th percentiles of monthly volatility in foot traffic across deciles. Establishments with lower monthly foot traffic show notably unstable foot-traffic growth patterns compared to those with higher foot traffic.

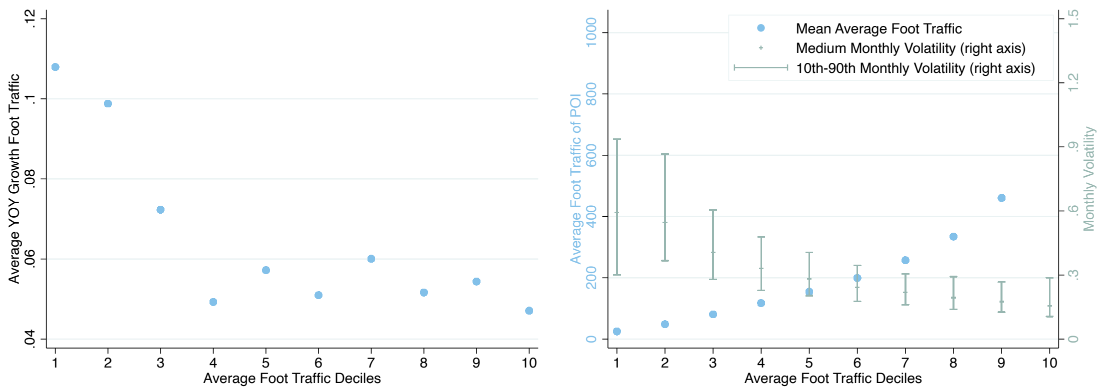


(a) Average YoY Growth

(b) Monthly Volatility

Figure C19: Growth Heterogeneity Among POIs By Deciles

## C.3 Monthly Volatility and YoY Growth Across Space and Sector

| | Monthly Volatility | | | | YoY Growth | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean | St. Dev | Min | Max | Mean | St. Dev | Min | Max |
| Overall | 0.34 | 0.21 | 0.01 | 2.40 | 0.06 | 0.20 | -0.66 | 0.99 |
| *By Borough* | | | | | | | | |
| Manhattan | 0.36 | 0.22 | 0.01 | 2.36 | 0.03 | 0.21 | -0.66 | 0.99 |
| Brooklyn | 0.32 | 0.20 | 0.01 | 1.69 | 0.12 | 0.20 | -0.54 | 0.99 |
| Bronx | 0.28 | 0.19 | 0.08 | 2.40 | 0.09 | 0.17 | -0.40 | 0.87 |
| *By Sector* | | | | | | | | |
| Retail | 0.36 | 0.23 | 0.01 | 2.36 | 0.08 | 0.21 | -0.66 | 0.99 |
| Service | 0.31 | 0.19 | 0.01 | 2.40 | 0.05 | 0.19 | -0.61 | 0.99 |
| Health | 0.49 | 0.26 | 0.11 | 1.96 | 0.11 | 0.22 | -0.41 | 0.82 |

Table C3: Monthly Volatility and Average Year-Over-Year Growth per POI

Table C3 provides descriptive statistics on monthly volatility and average year-over-year growth for establishments, categorized by different boroughs or sectors. The statistics suggest potential differing demand dynamics across locations and sectors. Specifically, Manhattan establishments show nearly zero average year-over-year growth but exhibit the highest volatility in monthly foot traffic. Conversely, establishments in the Bronx display faster average growth compared to the other boroughs, coupled with high volatility in their growth processes. Additionally, establishments in the health sector, on average, experience the highest growth rate and exhibit the highest monthly volatility compared to those in the other two sectors.

Table C4 presents the mean of monthly volatility and average year-over-year growth for establishments in each borough and each sector. We find that establishments in Brooklyn, across all sectors, exhibit the highest average year-over-year growth compared to other boroughs. Conversely, establishments in Manhattan, regardless of sector, consistently show the highest average monthly volatility. In all boroughs, estab-

| | Monthly Volatility | | | Average YoY Growth | | |
|---|---|---|---|---|---|---|
| | Retail | Service | Health | Retail | Service | Health |
| Manhattan | 0.42 | 0.32 | 0.53 | 0.03 | 0.02 | 0.09 |
| Brooklyn | 0.33 | 0.29 | 0.49 | 0.14 | 0.10 | 0.14 |
| Bronx | 0.29 | 0.25 | 0.45 | 0.10 | 0.07 | 0.12 |

Table C4: Monthly Volatility and Average Year-Over-Year Growth per POI Across Locations and Sectors

lishments in the health sector exhibit the highest average year-over-year growth and the greatest monthly volatility compared to establishments in other sectors. These findings align with the observations discussed in section 3.

## C.4 Other Characteristics of Establishments in Each Cluster

This section reports the top ten establishment categories within each cluster of the branded sample. The High-Traffic cluster is notably more concentrated in restaurants (47.8%) compared to the other two groups. This concentration may help explain why establishments in the high-traffic cluster consistently exhibit lower residual standard deviations than those in the Low-Traffic and Fast-Growing clusters.

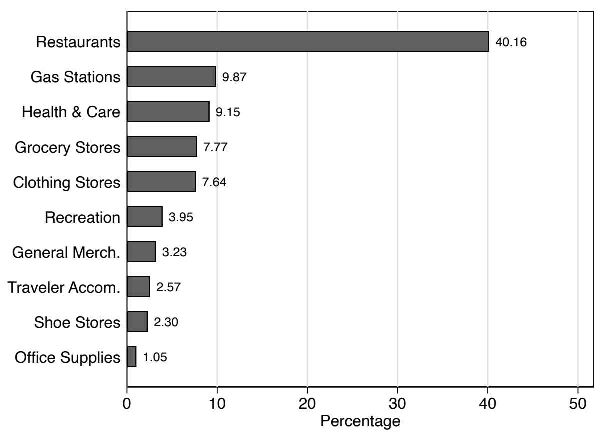


(a) Fast-Growing

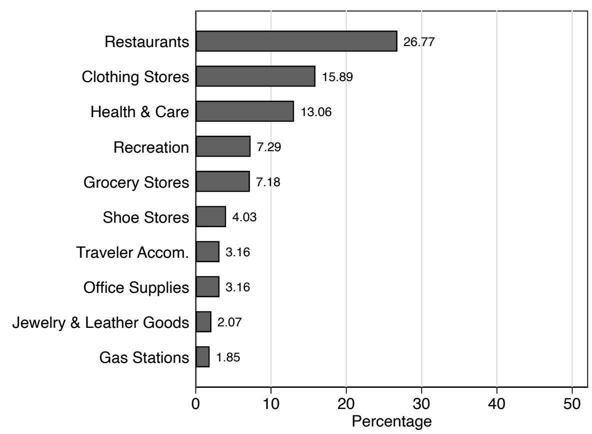


(b) Low-Traffic

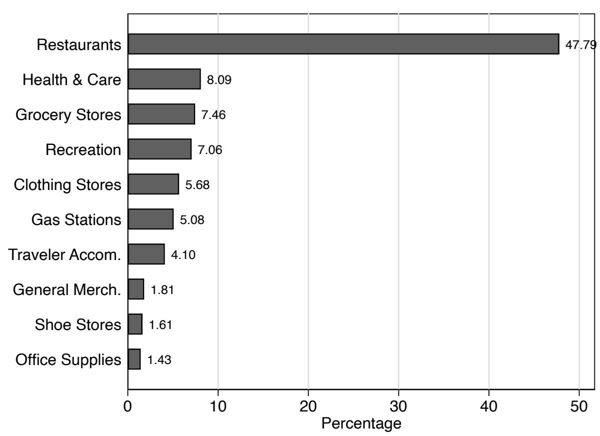


(c) High-Traffic

Figure C20: Top 10 Categories of Establishments

## C.5 AR(1) Demand Process

We begin by assuming a AR(1) process for the year-over-year demand growth rate with i.i.d errors for all branded establishments in each cluster $k$.

We applied the same method to form clusters, so the cluster groups remain the same as those shown in Table 2 in the main text. The estimates are based on equations (3) and (5), and we conducted a portmanteau test for autocorrelation of residuals, as introduced by Inoue and Solon (2006). The null hypothesis is that there is no first-order autocorrelation in the residuals. Since the p-value for all cluster groups is approximately 0, the null hypothesis is rejected, indicating the presence of autocorrelation in the residuals, which violates the i.i.d. error assumption. Thus, we incorporate a moving average process into the errors to account for the autocorrelation.

| Variable: - Cluster | Fast-Growing | Low-Traffic | High-Traffic | Pooled |
|---|---|---|---|---|
| *FT Characteristics* | | | | |
| $\overline{\Delta}_i$ | 0.289 | 0.073 | -0.035 | 0.065 |
| $\sigma_{\Delta_i}$ | 0.316 | 0.718 | 0.245 | 0.337 |
| $\overline{d}_i$ | 201.282 | 61.186 | 358.862 | 272.277 |
| *Sectoral Compositions:* | | | | |
| % in Retail | 0.492 | 0.534 | 0.393 | 0.440 |
| % in Service | 0.473 | 0.383 | 0.590 | 0.528 |
| *Location Compositions* | | | | |
| % in Manhattan | 0.421 | 0.681 | 0.587 | 0.559 |
| % in Brooklyn | 0.414 | 0.226 | 0.242 | 0.283 |
| *Regression Estimates* | | | | |
| $\rho$ | 0.151*** | -0.047** | 0.123*** | 0.084*** |
| | [0.016] | [0.019] | [0.008] | [0.008] |
| $\sigma_{u_i}$ | 0.141 | 0.237 | 0.098 | 0.181 |
| $\sigma_{\epsilon_{i,t}}$ | 0.393 | 0.521 | 0.216 | 0.342 |
| *Inoue and Solon (2006) Test:* | | | | |
| $P - value$ | 0.000 | 0.000 | 0.000 | 0.000 |
| *Fraction of Establishments* | 0.256 | 0.155 | 0.588 | |

*Note:* Standard errors are reported in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C5: Characteristics & Estimates of Clusters

# D Descriptive Statistics of the Full Sample

In this section, we present the descriptive statistics of the full sample for comparison with the branded establishments. Table D6 shows the summary statistics for adjusted foot traffic per establishment across years, locations, and sectors. The distribution of establishments by location mirrors that of the brand sample in Table 1, with over half situated in Manhattan, which records the highest average foot traffic. The full sample, however, has a higher proportion of health sector establishments (about 13% of the total) compared to the brand sample. As in the brand sample, establishments in the service sector have the highest average foot traffic, with the health sector the lowest.

| | # of POIs | Mean FT | St. Dev. FT |
|---|---|---|---|
| Overall | 44,011 | 171 | 300 |
| *By Year* | | | |
| 2018 | 44,011 | 172 | 306 |
| 2019 | 44,011 | 169 | 300 |
| *By Borough* | | | |
| Manhattan | 22,799 | 187 | 357 |
| Brooklyn | 15,670 | 142 | 205 |
| Bronx | 5,542 | 184 | 256 |
| *By Sector* | | | |
| Retail | 15,215 | 151 | 244 |
| Service | 23,016 | 203 | 344 |
| Health | 5,780 | 92 | 210 |

Table D6: Adjusted Foot Traffic (Full Sample)

| | Retail | Service | Health |
|---|---|---|---|
| Manhattan | 49%/33% | 57%/56% | 41%/11% |
| Brooklyn | 37%/36% | 33%/49% | 41%/15% |
| Bronx | 14%/39% | 10%/42% | 18%/19% |

Table D7: Distribution of POIs Across Locations and Sectors (%)

Table D7 presents the distribution of establishments by sector and location within the full sample. Within each sector, the largest proportion of establishments is located in Manhattan. Unlike the results in the brand sample (see Table D7), the majority of establishments across all boroughs in the full sample belong to the service sector. Moreover, establishments in the health sector represent a larger percentage in the full sample across all boroughs.