

Identifying Agglomeration Spillovers: Evidence from Grocery Store Openings*

August Chen[†], Franklin Qian[‡], Qianyang Zhang[§] & Xiang Zhang[¶]

March 2026

Abstract

We provide the first large-scale causal estimates of demand-side agglomeration spillovers in the non-tradable service sector, leveraging data on 851 grocery store openings across the United States between 2019 and 2022. We combine a convolutional neural network with propensity score matching to identify credible counterfactual opening sites and compare business outcomes between actual and counterfactual locations. Grocery store openings increase foot traffic to nearby businesses within 0.1 miles by 23.4 percent within 6–12 months, with spillovers decaying sharply beyond this radius. The effects are strongest for wholesale and retail establishments and hospitality services, and are substantially attenuated after COVID-19—from 42.7 to 14.6 percent—consistent with a trip-chaining mechanism. Grocery anchors also reshape the local business landscape: within 0.1 miles, establishment growth rates rise by 7.0 percentage points and employment growth rates by 8.8 percentage points, averaged over the first three post-entry years. Decomposing establishment growth rates, the effect is accounted for by higher entry and fewer exits.

*We would like to thank Brent Ambrose, David Brasington (discussant), Donald Davis, Jonathan Dingel, Jeffery Fisher, Jessie Handbury, Matthias Hoelzlein (discussant), Erik James (discussant), Niels Kuiper (discussant), Yuhei Miyauchi (discussant), Michael Pollmann, Steve Redding, Lindsay Relihan (discussant), Gregor Schubert, Jacob Sagi, Philip Ushchev (discussant), David Weinstein, Abdullah Yavas, and Tingyu Zhou (discussant) for helpful comments and suggestions. Shizhen Fan provided excellent research assistance. We also wish to thank the participants of the NBER Summer Institute (Real Estate), the Online Spatial and Urban Seminar, UEA-North America, UEA-Europe, the AREUEA-ASSA conference, the AREUEA national conference, the RERI conference, CICF, and the Eastern Finance Association annual meeting. We acknowledge financial support from the International Economics Section at Princeton University and the Real Estate Research Institute (RERI). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors. They do not necessarily reflect the views of RERI.

[†]Rice University, Jones Graduate School of Business. Email: aychen@rice.edu

[‡]UNC Kenan-Flagler Business School, Department of Finance. Email: franklin_qian@kenan-flagler.unc.edu.

[§]Amazon — Author affiliation reflects current position; this research was conducted prior to Qianyang Zhang’s employment at Amazon. Email: qianyangzz0704@gmail.com.

[¶]Princeton University, Department of Economics. Email: xiangzhang@princeton.edu.

1 Introduction

Grocery stores are among the most powerful anchors in American retail. In the United States, most shopping centers are anchored by at least one nationally recognized retailer that serves as the primary draw of customer traffic, and grocery-anchored centers have emerged as especially resilient assets—even as the broader retail sector has experienced widespread store closures since 2010 (Fung, 2020). The commercial real estate industry widely presumes that a new grocery anchor lifts demand for neighboring tenants, and the size of these spillovers matters for store location and tenant mix decisions. However, rigorous causal evidence on the magnitude of these spillovers remains scarce. How large are the demand externalities that grocery stores generate for nearby businesses? Do these externalities translate into real changes in local business formation and employment? This paper provides the first large-scale causal estimates of agglomeration spillovers from grocery store openings in the non-tradable service sector, leveraging a novel identification strategy that combines deep learning with propensity score matching.

Using a comprehensive sample of 851 grocery store openings across the United States between 2019 and 2022, we estimate the causal effects of grocery store entry on two classes of outcomes: (i) demand spillovers to nearby incumbent businesses, measured by foot traffic, and (ii) the dynamics of surrounding businesses, measured by net growth in establishment counts and employment. To isolate causal effects, we develop a two-stage identification strategy. In the first stage, a generative adversarial network (GAN)-based convolutional neural network (CNN) screens continuous geographic space to identify a pool of candidate counterfactual locations that resemble actual opening sites in their demographic and business-environment characteristics. In the second stage, we estimate propensity scores using a parsimonious set of economically interpretable covariates and match each real opening to its nearest-score counterfactual from the CNN-derived candidate pool. This approach combines the pattern-recognition strength of deep learning for spatial screening with the transparency of traditional propensity score methods, addressing a key challenge in spatial causal inference: finding credible counterfactual locations in continuous space where the set of potential controls is vast and high-dimensional.

We first validate the identification strategy by confirming substantial overlap in propensity score distributions between real openings and matched counterfactual sites, balanced covariates across the two groups, and a sharp 404 percent increase in foot traffic at the grocery store locations themselves within 6–12 months of opening—lending credence to both the accuracy of recorded opening dates and the relevance of the treatment.

Turning to demand spillovers, we find that grocery store openings generate economically large and highly localized positive externalities on nearby businesses. Average foot traffic to businesses within 0.1 miles of a new grocery store increases by 23.4 percent in the 6–12 months following the opening, relative to businesses surrounding matched counterfactual sites. These spillovers decay sharply with distance: beyond 0.1 miles, the estimated effects are statistically indistinguishable from zero. The demand spillovers exhibit substantial heterogeneity across business categories. Wholesale and retail stores (excluding grocery stores) experience an average increase of 30 percent

in foot traffic, and the strongest synergies are observed between grocery stores and wholesale and retail establishments as well as hospitality services for accommodations, eating, and drinking—categories that share the most customer overlap with grocery shoppers. By contrast, medical, welfare, and healthcare services exhibit a statistically significant but economically modest response of approximately 10 percent—20 percentage points smaller than the wholesale and retail effect—consistent with weaker demand complementarity with grocery trips. Notably, we find no evidence of negative competitive effects on nearby incumbent grocery stores on average.

The spillovers also vary meaningfully by the type of entering grocery store. Openings of national and regional chains (e.g., Whole Foods Market, Trader Joe’s) generate the strongest effects, with a 31 percent increase in foot traffic to surrounding businesses, while big-box retailers and discounters generate more modest spillovers of 16 and 17 percent, respectively. This pattern is consistent with national chains attracting higher-frequency, multi-purpose shopping trips that spill over to adjacent businesses. Strikingly, we find that demand spillovers are substantially attenuated after the onset of the COVID-19 pandemic: pre-COVID openings increase nearby foot traffic by 42.7 percent, compared with only 14.6 percent for post-COVID openings. A back-of-the-envelope calculation suggests that roughly 90 percent of this decline is attributable to lower foot traffic at the grocery stores themselves. This finding is consistent with the trip-chaining mechanism ([Miyauchi et al., 2025](#)): as remote work reduced commuting and allowed consumers to decouple routine grocery runs from discretionary shopping stops ([Qian and Su, 2024](#)), the incidental foot traffic that spills over to neighboring businesses diminished proportionally.

Beyond demand spillovers to incumbents, we document that grocery store openings reshape the surrounding business landscape. Averaging over the first three post-entry years, establishment growth rises by 7.0 percentage points and employment growth rises by 8.8 percentage points within 0–0.1 miles of an opening site, relative to matched counterfactual sites. These effects attenuate quickly: in the 0.1–0.2 mile ring, the corresponding increases are 2.0 and 3.0 percentage points, and beyond 0.2 miles the estimates are indistinguishable from zero. The growth rate effects are largest in the first post-entry year—12.4 percentage points for establishments and 15.8 percentage points for employment—and return to near zero in subsequent years. Because these coefficients measure growth rate differentials between treated and control sites, a return to zero does not imply that business activity declines; rather, it indicates that the treated group’s growth rate converges back to the control group’s rate, so that the cumulative gains in establishment counts and employment continue to build over the three-year post-entry window. Decomposing establishment growth rates, the 7.0 percentage point average effect is accounted for by both higher entry (4.2 percentage points) and fewer exits (2.9 percentage points)—the only margins that operate when the outcome is a binary active-status indicator. Industry-level estimates indicate that the localized gains are concentrated in complementary sectors with strong linkages to grocery-driven demand—retail-adjacent activities, restaurants, and selected local services—underscoring that grocery anchors catalyze localized business formation in the immediate vicinity of the new store. These business dynamics results reveal a second channel through which grocery anchors affect the local economy:

beyond raising foot traffic at incumbent businesses, the demand spillovers enhance the commercial viability of nearby locations, attracting new entrants and sustaining marginal incumbents that might otherwise exit.

Our paper contributes to several strands of literature. First, we contribute to a large body of research on the economics of agglomeration. A central finding in this literature is that the spatial concentration of economic activity raises the productivity of firms in manufacturing and tradable services, through channels such as reduced transportation costs, knowledge spillovers, labor market pooling, and improved firm-worker matching.¹ Recent work by [Baum-Snow et al. \(2024\)](#) documents that these productivity spillovers operate at remarkably fine spatial scales—within 75 meters—in the high-skilled services sector, driven primarily by knowledge transfer between nearby firms. Within the non-tradable service sector, the agglomeration literature has largely focused on how shopping center owners internalize anchor-tenant externalities through differential rental contracts ([Pashigian and Gould, 1998](#); [Gould et al., 2005](#); [Konishi and Sandfort, 2003](#); [Wheaton, 2000](#)).² A separate literature documents the negative externalities of firm closures and bankruptcies in this sector ([Shoag and Veuger, 2018](#); [Bernstein et al., 2019](#); [Benmelech et al., 2019](#); [Knight, 2022](#); [Billings et al., 2024](#)), and a body of work examines the mixed effects of big-box retail entry on local labor markets.³ Despite this extensive literature, causal estimates of *positive* demand-side agglomeration spillovers in the non-tradable sector remain rare.⁴ We contribute by providing the first large-scale causal estimates of demand spillovers from business openings in this sector, quantifying both their magnitude and spatial reach at a fine geographic resolution. Our results show that grocery anchors generate foot traffic externalities that decay sharply beyond 0.1 miles—a level of spatial precision that is invisible in the county-level analyses typical of the big-box literature. Moreover, our decomposition of business dynamics into extensive and intensive margins reveals that grocery anchors not only attract more new entrants but also reduce exits of incumbent businesses, a pattern distinct from the negative spillovers documented in the closure literature.

Second, we contribute to the literature on the demand mechanisms underlying agglomeration in the non-tradable sector. Within retail settings, shared customer traffic is widely considered the primary driver of spatial spillovers, as clustering reduces consumer search costs and enables trip chaining—the practice of combining multiple shopping stops into a single outing ([Wolinsky, 1983](#); [Pashigian and Gould, 1998](#); [Gould et al., 2005](#)). Recent theoretical and empirical advances have formalized trip chaining as a key microfoundation for consumption externalities across locations ([Miyauchi et al., 2025](#); [Oh and Seo, 2023](#); [Relihan, 2022](#); [Su, 2022](#)).⁵ We provide direct empirical

¹For recent surveys, see [Duranton and Puga \(2004\)](#), [Glaeser and Gottlieb \(2009\)](#), [Moretti \(2011\)](#), and [Rosenthal and Strange \(2003\)](#). Foundational contributions include [Krugman \(1991\)](#). See also [Greenstone et al. \(2010\)](#), [Ellison and Glaeser \(1997\)](#), [Ellison et al. \(2010\)](#), [Glaeser et al. \(1992\)](#), [Henderson et al. \(1995\)](#), and [Henderson \(2003\)](#).

²See also [Benjamin et al. \(1992\)](#), [Gatzlaff et al. \(1994\)](#), [Liu and Liu \(2013\)](#), and [Agrawal and Cockburn \(2003\)](#).

³See, for example, [Basker \(2005\)](#), [Neumark et al. \(2008\)](#), [Jia \(2008\)](#), [Merriman et al. \(2012\)](#), [Ellickson and Grieco \(2013\)](#), [Bertrand and Kramarz \(2002\)](#), and [Sadun \(2015\)](#).

⁴Other related papers include [Zhou and Clapp \(2015\)](#), [Liu et al. \(2018\)](#), [Clapp et al. \(2019\)](#), [Rosenthal and Strange \(2020\)](#), [Liu et al. \(2020\)](#), [Kuiper et al. \(2021\)](#), [Liu et al. \(2026\)](#), and [Hu et al. \(2025\)](#).

⁵Related work by [Leung and Li \(2022\)](#) shows that store format shapes consumer shopping behavior and welfare, and [Qian and Tan \(2022\)](#) and [Rosenthal and Urrego \(2023\)](#) document additional externalities of spatial concentration

evidence that trip chaining operates as a first-order agglomeration mechanism (formalized in the conceptual framework of Appendix D, which derives two empirical predictions linking variety gains and trip-bundling savings to the spatial and cross-category patterns we document). By exploiting the COVID-19 pandemic as a natural experiment that disrupted multi-purpose shopping behavior, we show that the sharp attenuation of spillovers post-COVID—from a 42.7 to a 14.6 percent increase in nearby foot traffic—is quantitatively consistent with reduced trip bundling. This finding offers a clean test of the trip-chaining channel that is difficult to obtain in cross-sectional settings.

Third, our paper contributes to the growing literature on machine learning and artificial intelligence in economics and real estate. A rapidly expanding body of work demonstrates that machine learning tools can enhance both prediction and causal inference in economic settings (Mullainathan and Spiess, 2017; Athey and Imbens, 2019).⁶ In real estate and urban economics, ML methods are increasingly applied to property valuation, market analysis, and the modeling of commercial real estate transactions (Kojen et al., 2025).⁷ Our paper bridges the ML prediction literature and the causal inference tradition: we demonstrate that convolutional neural networks, originally developed for image recognition, can serve as a powerful spatial screening tool within a causal inference pipeline. The CNN identifies candidate counterfactual locations from continuous geographic space—a task where the high dimensionality of spatial features makes traditional methods impractical—and the subsequent propensity score stage ensures that the final matched pairs are balanced on economically interpretable covariates. This two-stage approach may generalize to other spatial treatment-effect settings where credible counterfactuals must be selected from a vast set of potential control locations.

Fourth, we contribute to the literature on causal inference in spatial settings. Traditional spatial treatment designs compare outcomes in an inner ring to those in an outer ring around each treatment site (Currie et al., 2015; Ellickson and Grieco, 2013), but distance to a treatment location often correlates with both observed and unobserved confounders, which undermines identification. Recent work has developed non-parametric (Diamond and McQuade, 2019) and semi-parametric (Butts, 2023) approaches to spatial spillover estimation. We extend Pollmann (2020)’s GAN-based CNN approach into a scalable identification strategy that selects counterfactual locations closely matched to actual openings. Relative to Pollmann (2020), our contributions are fourfold: we expand the CNN’s input set to include richer multi-scale features encompassing demographic, economic, and business-environment variables; we improve predictive accuracy through a heatmap-labeling loss component⁸ that conveys distance-weighted feedback during training; we extend the geographic scope from a single metropolitan region to grocery store openings across the entire United States; and we leverage the sharp timing of store openings as quasi-experimental shocks in an event-study framework, enabling pre-trend validation and precise treatment-effect estimation.

after business entries.

⁶See also Kleinberg et al. (2018) on how ML predictions improve decision-making, and Fuster et al. (2022) on ML in credit markets.

⁷See, for example, Shen and Ross (2021), Lorenz et al. (2023), and Wan and Lindenthal (2023) for applications of ML methods in real estate.

⁸We refer to our method of encoding spatial correctness via heatmaps in the loss function as *heatmap labeling*.

The remainder of the paper is organized as follows. Section 2 discusses the data used for the analysis. Section 3 outlines the empirical design. Section 4 reports the main results. Section 5 concludes.

2 Data

2.1 Sample of Grocery Store Openings

We construct our dataset of grocery store openings using two complementary data sources. Our primary source consists of store opening information collected directly from the official websites of major grocery chains. We begin with a list of leading U.S. grocery chains compiled by Gourmet Pro, a consulting firm specializing in the grocery retail industry. Gourmet Pro ranks chains based on annual revenue, number of stores, geographic coverage, customer base, brand strength, and operational scale. For each chain on this list, we manually collect information on store opening locations, addresses, and opening dates from the firm’s official website. We are able to obtain the address and opening month for all newly opened stores between 2019 and 2022 for 11 chains. These chains collectively generate more than \$980 billion in annual revenue and operate over 15,000 stores nationwide.⁹

To broaden coverage beyond the largest national chains, we supplement this dataset with store opening information from Chain Store Guide (CSG), a professional retail data vendor. CSG provides store names and addresses for additional openings, particularly among smaller and regional chains. Because opening dates reported in CSG are often imprecise, we manually verify the opening month using local news sources and press releases and retain only stores for which the opening month can be confirmed.

We geocode store addresses and match them to commercial points of interest (POIs) from SafeGraph to measure foot traffic, defined as the monthly number of visits to each store. We further restrict the sample to stores for which surrounding businesses can be identified within the predefined radius used in the analysis. After applying these restrictions, the final sample consists of 851 grocery store openings.

Figure 1 displays the geographic distribution of store openings in the sample. We classify grocery stores into three categories: (1) big-box retailers, accounting for 14.6 percent of openings; (2) discounters, accounting for 46.2 percent of openings; and (3) national and regional grocery chains, accounting for 39.2 percent of openings. Appendix Table A1 lists the individual grocery chains, their classification into these three store types, and the number of openings for each chain.

⁹For the full list of chains, see <https://www.gourmetpro.co/blog/largest-grocery-chains-usa>. We exclude pharmacy chains and dollar/variety stores (e.g., CVS, Walgreens, and Dollar General). The chains for which detailed opening information is available from official websites include Walmart, Costco, Safeway, Food Lion (Ahold Delhaize USA), Publix, ALDI, Target, and Whole Foods.

2.2 Construction of Outcome Variables

2.2.1 Foot Traffic to Business Establishments

We obtain the monthly number of visits to commercial POIs from SafeGraph, accessed via the Dewey Platform. SafeGraph covers the vast majority of POIs in the U.S. and compiles foot traffic data from multiple sources, including mobile devices, WiFi connections, and sensors. In addition, SafeGraph provides POI characteristics such as their names, addresses, and industry classifications. We classify commercial POIs into six categories: 1) Grocery and Competitors; 2) Wholesale and Retail; 3) Accommodation, Eating, and Drinking; 4) Medical, Welfare, and Healthcare; 5) Finance, Real Estate, Communication, and Professional Services; 6) Other Services, following [Miyauchi et al. \(2025\)](#).¹⁰

2.2.2 Counts of Business Establishments and Employment

We obtain counts of business establishments and their annual employment from Data Axle (formerly Infogroup). The database covers 226 million verified U.S. businesses dating back to 1997, providing the establishment-level names, street addresses, longitudes, and latitudes, and employee counts. The annual snapshots of businesses also allow us to track the entries and exits of businesses over the years. We use the number and employee counts of businesses that are located within 0.5 miles of the grocery store openings and their matched counterfactual sites from 2014 to 2024 in our analysis of business dynamics.

3 Empirical Approach

After a grocery store opening, changes in foot traffic to nearby businesses may reflect three conceptually distinct channels: (1) *Localized spillovers from entry*—the grocery store as an anchor not only brings new foot traffic to itself, but also draws customer traffic to nearby businesses through increased salience, reduced search costs, and trip chaining. (2) *Treatment-induced broad general-equilibrium (GE) effects*—the entry may generate neighborhood-wide changes, such as increased residential demand or business investment, that extend beyond the immediate vicinity of the opening. (3) *Area-wide trends independent of the entry*—the entire area surrounding the opening could be trending up in business densities and customer demand due to local economic conditions unrelated to the grocery store. In the third channel, the increase in foot traffic to nearby businesses would have occurred in the absence of the grocery store entry, as the entire area is on an upward growth trajectory. A key challenge for identification is to separate the first two channels—which are causal effects of the opening—from the third channel, which is a confounder.

¹⁰Wholesale and retail stores are identified by 2-digit NAICS codes 42, 44, 45; accommodations, eating, and drinking places are identified by 2-digit NAICS code 72; finance, real estate, communication, and professional services are identified by 2-digit NAICS codes 52, 53, 54, 55, 56; medical, welfare and healthcare stores are identified by 2-digit NAICS code 62; and other services include 2-digit NAICS codes 51, 61, 71, 81, 92. We summarize the definitions of business categories in Appendix Table A2.

To identify the treatment effects of grocery store openings on foot traffic to surrounding businesses, a simple strategy would be to compare foot traffic to businesses close to the opening with that to businesses further away via a spatial difference-in-difference design. However, the key confounding issue is that businesses close to the opening could be unobservably different from those further away, and their distance to the opening is correlated with their likelihood of attracting a grocery store nearby. For example, businesses that are closer to the opening could be in locations where the pre-existing densities of businesses are already high, and the options for grocery shopping are limited, so they are more likely to attract a grocery store opening nearby.

To overcome this identification challenge, ideally, we would like to compare the outcomes of businesses around a grocery store opening with those of businesses around an alternative location that was equally desirable for the opening but did not receive one. However, the desirability of a location for a grocery store depends on a high-dimensional set of covariates—neighborhood demographics, business composition, traffic access, and competitive proximity—that are difficult to summarize with standard tabulated matching covariates.

Our solution combines two complementary tools. First, we train a convolutional neural network (CNN) to read the spatial layout of demographics and businesses around a site—much as a retailer’s site-selection team would evaluate a map—and use it to identify a pool of candidate counterfactual sites in the vicinity of each actual opening. Second, we refine this pool via propensity-score matching on interpretable economic covariates, ensuring transparent covariate balance. We then compare outcomes of businesses at equal distances from real and matched counterfactual sites in an event-study framework. We outline each step of this approach in detail below.

3.1 Predict Counterfactual Locations for Grocery Store Openings

The core of our predictive model is a GAN-based CNN following [Pollmann \(2020\)](#). We extend his approach in three directions motivated by our empirical setting: (i) grocery-store location decisions depend on a wider set of neighborhood characteristics than the restaurant setting in [Pollmann \(2020\)](#), so we enrich the input features to include demographics, purchasing power, traffic access, and detailed business composition;¹¹ (ii) because our counterfactual sites must be precisely localized to define concentric rings for the event study, we introduce heatmap labeling to improve the CNN’s spatial precision; and (iii) we demonstrate that the approach generalizes across a broad set of U.S. metro areas and geographies, rather than a single city.

Our model construction proceeds in three main stages. First, we select relevant variables and transform the raw data into a spatial dataset with economically meaningful features suitable for CNN training. Second, we identify a pool of potential counterfactual sites by augmenting the CNN with a generative adversarial network (GAN), which allows the model to efficiently evaluate and select candidate locations similar to actual grocery store openings in demographic and business characteristics. Third, we address the rare-treatment problem caused by the large candidate pool by

¹¹The features [Pollmann \(2020\)](#) uses are the location of surrounding grocery stores, restaurants, and a combination of all other kinds of business.

applying a two-stage propensity score matching procedure. This approach also reduces dependence on hyperparameter tuning and improves economic interpretability. Collectively, these three stages provide a robust control group, enabling credible causal inference regarding the impact of grocery store openings.

3.1.1 Input Structure and Sample Preparation for CNN

To capture grocery-store location decisions, we assemble 22 variables spanning four dimensions in the literature (Krishna, 2021; Waters, 2021): neighborhood demographics, site accessibility and traffic, customer purchasing power, surrounding business composition and density, and competitive proximity. For variables requiring mean statistics—such as population and per-capita income—we draw on the most recent American Community Survey (ACS) five-year estimates for Census Block Groups, using data from the years preceding each opening. To calculate the incumbent business density before the openings, we employ SafeGraph’s Points of Interest (POI) dataset, which provides detailed records of business establishments. The CNN model uses both the locations and counts of business POIs as inputs. Specifically, their latitude and longitude determine their positions in the input tensors, while their counts define the corresponding values in the tensors. We aggregate these data in semi-annual intervals and, for each half-year period, count only those POIs that commenced operations at least one month before the actual opening. We present the variables used as input features to the CNN model in Table 1.

To construct the input tensors for the CNN model, we discretize the continuous geographic area around each opening into a 10×10 grid of 0.05-mile cells, a resolution chosen to balance predictive accuracy against computational efficiency: larger cells decrease spatial precision, while smaller cells impose high computational costs with marginal accuracy gains. Within each cell, we summarize continuous variables (e.g., income, population share) by their mean and count-based variables (e.g., business density) by their totals. This procedure yields a $23 \times 10 \times 10$ tensor that encodes the spatial distribution of our covariates and serves directly as input for our CNN.

Because our sample of 851 grocery-store openings is small by deep-learning standards, we apply standard data-augmentation techniques—random translations, rotations, and reflections—to expand the effective training set.¹² Since all variables are computed spatially relative to the opening, these distance-preserving transformations leave the variable values unchanged, so each augmented input retains all authentic spatial information.

After preprocessing, we split the data into an 80 percent training subset and a 20 percent validation subset. Once the CNN demonstrates stable out-of-sample performance, we retrain it on the full dataset to obtain our final network. We then deploy the trained network to evaluate a large set of randomly sampled candidate sites around each actual store opening, using the CNN to assess the suitability of each location. This step is intended to screen continuous geographic space and identify plausible counterfactual candidates with similar high-dimensional local characteristics.

¹²Deep-learning models typically require tens of thousands of labeled examples for robust generalization (e.g., Krizhevsky and Hinton, 2009; Deng et al., 2009). Our augmentation techniques are visualized in Appendix Figure B1a.

From this pool, we retain the top-ranked candidate sites (as determined by the CNN-generated scores) as the basis for subsequent propensity-score matching.

3.1.2 GAN-Based CNN Architecture

The CNN’s task is to learn what a grocery-desirable location looks like from its spatial environment—and then identify other locations that look equally desirable but lack a grocery store. To train the network for this task, we construct three types of grid inputs (see Appendix Fig. B1b): (Type I) grids where we take a real grocery opening and erase it from the map, so the surrounding environment still “looks like” a grocery location but the store itself is missing; (Type II) grids where the real opening is left intact; and (Type III) random locations with no plausible openings. The CNN must learn to distinguish these three cases. For Type I grids—our constructed counterfactuals—the network is further trained to pinpoint the exact cell where the erased opening was located.

We enhance this architecture in two ways. First, we embed adversarial (GAN) training principles (Goodfellow et al., 2020) to sharpen the network’s discrimination between Type I (grocery “missing”) and Type II (grocery present) grids. The adversarial training forces the CNN to become increasingly discriminating between locations where a grocery store is missing versus locations where one is already present or implausible, improving its ability to detect subtle spatial signatures of grocery-desirable sites. Second, we introduce heatmap labeling for the localization stage: rather than binary correct/incorrect feedback for each predicted cell, the loss function assigns a continuous penalty based on distance from the true opening location, providing graduated spatial feedback that improves localization precision (see Appendix Figure B2).

After final training, the GAN-based CNN classifies tensors into Type I, II, and III with accuracies of 99.7%, 99.4%, and 98.8%, respectively—meaning it misclassifies fewer than 12 in every 1,000 samples. When we focus on Type I tensors, the network pinpoints the exact missing-opening cell 99.2% of the time, and if we allow a one-cell spatial tolerance (i.e., any of the eight neighboring cells), accuracy rises to 99.7%.

3.1.3 Propensity-Score Estimation and Matching

Using the trained CNN model, we generate a pool of candidate counterfactual sites around each actual opening. We then refine this pool via propensity-score matching on our predetermined economic covariates.

While the CNN excels at spatial pattern recognition, it does not guarantee balance on specific economic covariates. We therefore apply propensity-score matching to the CNN-screened candidate pool, ensuring transparent covariate balance on interpretable economic indicators.

The matching covariates include demographic features—population, age and racial composition, educational attainment, household income, and unemployment—computed as population-weighted spatial averages of Census Block Group (CBG) attributes over a five-mile local demographic aggregation window around each site.¹³ For the business-density covariates, we compute five industry-

¹³Specifically, each CBG’s contribution is weighted by the product of its population and the share of its area that

specific densities within a 0.5-mile radius of each site by counting SafeGraph POIs that commenced operations at least one month before the opening. We divide each category’s count by the area of the circle to obtain densities for “Wholesale and Retail,” “Finance, Real Estate, Communication, and Professional,” “Accommodations, Eating, and Drinking,” “Medical, Welfare, and Healthcare,” and “Grocery and General Merchandise Retailers.” We treat the first four categories as measures of complementary business presence and the fifth as a measure of local competition, thereby capturing both supportive and competitive commercial environments in our matching covariates. The detailed NAICS code definitions for these five industry categories are provided in Appendix Table A2.

The CNN model’s ability to generate millions of candidate counterfactual locations creates a severe imbalance between the few treated sites (the actual openings) and the vast pool of potential controls, risking a rare-treatment violation of the overlap assumption. To mitigate this issue, we implement a two-stage propensity-score matching strategy that follows a coarse-to-fine procedure. In the first stage, we exclude controls whose estimated propensity scores fall outside the common-support region of the treated units and pair each treated site with its ten nearest counterfactuals (1-to-10 matching) based on these initial scores. In the second stage, we re-estimate propensity scores on the pruned sample and conduct nearest-neighbor matching with up to three control sites per treated site. Specifically, each real opening is matched to its up-to-three nearest counterfactual sites in propensity score, within a caliper equal to 0.2 times the standard deviation of the logit propensity score, as recommended by Austin (2011); if fewer than three counterfactuals fall within the caliper, the treated site is matched to one or two controls. This procedure follows the propensity score matching framework of Rosenbaum and Rubin (1985). We also impose a geographic distance constraint, requiring matched counterfactual sites to be located at least 0.2 miles and no more than 15 miles from the corresponding real opening. The minimum distance ensures that the counterfactual site does not overlap spatially with the treatment site’s immediate vicinity, while the maximum distance ensures that matched pairs share similar unobserved local market conditions. In practice, most matched pairs for urban openings are separated by fewer than 5 miles; larger separations arise primarily in suburban and rural areas where grocery catchment areas are wider.

We validate the matching quality—including propensity-score overlap, covariate balance, and the realization rates of counterfactual sites as actual future openings—in Section 4.

3.2 Main Specification for Treatment Effects of Grocery Store Openings on Surrounding Businesses

Figure B3 presents two illustrative real-counterfactual (CF) pairs—one pre-COVID and one post-COVID—to show how CF sites compare to their matched real openings. In the first pair, a Trader Joe’s opened in September 2018 at One Loudoun in Ashburn, VA; the matched CF site lies approximately 5 miles east at Sugarland Crossing. In the second, a small-format Target opened in October 2020 at Shepherd Square in Houston, TX; the matched CF site is approximately 12 miles falls within the five-mile radius. The formal weighting equations are provided in Appendix C.

southwest in Stafford. In both cases, the matched CF sites exhibit comparable business density, road connectivity, and retail composition within a 0.5-mile radius. Notably, both CF sites were subsequently developed into grocery-anchored retail locations—the Ashburn CF site became a Lidl in 2021, and the Houston CF site became a Costco Business Center in 2021—providing independent face validity for the CNN’s ability to identify locations that are genuinely attractive to grocery retailers.¹⁴

Taking the Trader Joe’s opening as an example to illustrate our identification strategy, Trader Joe’s could have chosen the counterfactual site, which has a similar desirability to open its business. For idiosyncratic reasons of the owner of a particular grocery store that are not systematically correlated with the factors we account for in our CNN-pscore matching algorithm, the real site is chosen over the counterfactual site. In other words, these idiosyncratic factors that matter for the choice of a particular grocery contribute to the quasi-random variation we leverage for identification. However, such idiosyncrasies could not be correlated with any remaining systematic unobservables that matter for grocery stores’ location choices across our sample, that are not already accounted for in our predictive model. In the absence of the grocery store entry, businesses surrounding both the real and counterfactual sites would evolve on similar trajectories. After the grocery store opens, it becomes an anchor for shopping malls to draw customer traffic not only to itself but also to nearby stores. Relative to businesses around the counterfactual site, we expect foot traffic to businesses around the real opening to increase due to economies of agglomeration, rather than reasons related to local economic conditions alone. Based on this intuition, our identification strategy is to compare the outcomes of businesses that are at identical distances from the real and counterfactual opening sites, respectively.

To formalize what our design identifies, it is useful to decompose the outcome for business j in ring n around candidate site s associated with opening event i in period t :

$$Y_{i,s,n,t} = \underbrace{TE_{i,s,n,t}^{\text{local}}}_{\text{localized spillovers from entry}} + \underbrace{GE_{i,t}^{\text{treat}}(s)}_{\text{broad GE caused by entry}} + \underbrace{GE_{i,t}^0(s)}_{\text{area-wide shocks independent of entry}} + Shock_{i,s,t} + \varepsilon_{i,s,n,t}. \quad (1)$$

Here $GE_{i,t}^0(s)$ captures neighborhood-wide demand shocks or trends that are *independent* of the grocery entry—for example, a local boom or redevelopment that would occur even absent the opening. These are the confounders discussed above. $GE_{i,t}^{\text{treat}}(s)$ denotes treatment-induced general-equilibrium spillovers that may be spatially broad—for example, the opening may raise the profile of the wider neighborhood, attracting new residents or businesses to the area beyond the immediate vicinity of the store.

¹⁴A potential concern is that CF sites subsequently realized as actual openings could contaminate our control group if the later entry generates spillovers to surrounding businesses. We retain all CF sites in our main analysis because this concern is quantitatively limited in our setting. Our baseline treatment radius is 0.1 miles, and fewer than 3% of all realized openings occur within 0.2 miles of a previously matched CF site over the full sample period (2019–2022). In addition, most CF realizations occur more than one year after the corresponding real opening, placing them well outside our primary post-treatment window of 12 months. These spatial and temporal separations make it unlikely that later CF realizations materially affect our estimates.

Our real-versus-counterfactual comparison yields

$$Y_{i,s_i^*,n,t} - Y_{i,s_i^{CF},n,t} = (TE_{i,s_i^*,n,t}^{\text{local}} - TE_{i,s_i^{CF},n,t}^{\text{local}}) + (GE_{i,t}^{\text{treat}}(s_i^*) - GE_{i,t}^{\text{treat}}(s_i^{CF})) + (GE_{i,t}^0(s_i^*) - GE_{i,t}^0(s_i^{CF})), \quad (2)$$

where we suppress site-specific shocks for brevity. Two conditions determine what this difference identifies.

First, the confounding area-wide trends $GE_{i,t}^0$ must be common across the matched sites: $GE_{i,t}^0(s_i^*) \approx GE_{i,t}^0(s_i^{CF})$. Our CNN-propensity-score matching procedure is designed to ensure this condition holds, so that the third term in equation (2) vanishes. This is the core identification requirement.

Second, the treatment-induced broad GE component $GE_{i,t}^{\text{treat}}$ does *not* need to be zero for identification, but its spatial reach determines the interpretation of our estimand. If $GE_{i,t}^{\text{treat}}$ is common across the relevant local market—affecting businesses around both the real and counterfactual sites equally—then it also differences out, and our estimates capture only the localized spillovers TE^{local} . If instead the treatment-induced GE effects are spatially uneven, some portion remains in the estimand: our estimates then reflect localized spillovers plus any differential broad GE between the real and counterfactual locations. In either case, our design does not capture the *total* general-equilibrium impact of the opening, because the common component of $GE_{i,t}^{\text{treat}}$ is absorbed by the differencing. The estimates should therefore be read as a lower bound on the full causal effect of entry, obtained without imposing structural assumptions on the magnitude or spatial distribution of the broad GE channel. In the event-study specification below, the coefficient $\beta_{\tau,n}$ captures the average of TE^{local} across openings at distance band n and event time τ , plus any systematic differential in GE^{treat} between real and counterfactual sites.

To execute our identification strategy, for each grocery store opening, we define a case by pairing it with its matched counterfactual (CF) opening sites. We divide the businesses surrounding the real and CF opening sites into successive concentric rings, respectively. We then compare the outcomes of businesses located in the rings that are equally distant from the real and CF sites. The implicit assumption is that not only are real and CF sites directly comparable in their surrounding characteristics conducive to grocery store openings, but the businesses that are equidistant from the opening sites are also comparable, which we validate later, hence those near the CF sites would serve as a valid control group. This assumption is reasonable given that the effects of the openings of grocery stores are highly localized, as we will show later. We hence only examine businesses that are mostly within 0.1 miles from the real and CF sites, which ensures the surrounding environment would not differ very much.

To estimate the effects of grocery store openings on the outcomes of neighboring businesses, we use the following event study design as the baseline specification of our analysis:

$$Y_{ijnt} = \sum_{\tau \in [-s_1, s_2]} \beta_{\tau,n} \text{Treat}_{ij} \times D_{\tau,ijnt} + \alpha_{ij} + \phi_{inkt} + \varepsilon_{ijnt}, \quad n = 1, \dots, N \quad (3)$$

Here, Y_{ijnt} is the outcome of the business of interest j in distance band n associated with the

case of grocery i 's opening in time period t . For example, it could be the log of the monthly number of visits. The dummy variable Treat_{ij} takes the value of one if business j associated with the opening case i is next to the real opening and takes the value of zero if it is next to the CF site matched. The variable $D_{\tau,ijnt}$ denotes a dummy equal to one if the outcome Y_{ijnt} is observed in τ time periods relative to the opening of the grocery store i , where τ goes from s_1 periods before the opening to s_2 periods after the opening. In our baseline specification, we set $s_1 = 4$ and $s_2 = 12$, tracking outcomes from 4 months before the opening to 12 months after. α_{ij} denotes the opening case by business fixed effect that controls for time-invariant characteristics for each business j associated with opening case i . ϕ_{inkt} denotes the opening case by ring n by business category by time-period fixed effect. It allows for a flexible time trend for each business category in each local area that contains a pair of real and CF sites. Adding this term ensures that we only compare businesses belonging to the same category near the real and CF sites within the same local area defined by an (opening case, ring) to remove potential composition biases of businesses across geographic locations.¹⁵ Empirically, we can run a separate regression comparing rings that are equidistant from the real and CF sites for successive rings $n = 1, \dots, N$. Our coefficient of interest, which quantifies the average effects of an opening on nearby businesses τ time periods relative to the opening, is denoted by $\beta_{\tau,n}$. We normalize the coefficient of one month before the opening, $\beta_{-1,n}$, to zero. Standard errors are clustered at the real or CF site level.

Appendix D formalizes the demand-side mechanism underlying these regressions. In the model, a grocery opening affects visits to a nearby POI (k, j) through two channels: a *variety channel* that lowers the category-level price index P_{grocery} by adding a new consumption option, and a *trip-bundling channel* that lowers the effective travel cost $\bar{t}_{kj}$ for nearby businesses across all categories. Empirical Prediction I maps to the spatial decay pattern we test in Section 4: closer POIs experience larger visit increases because trip bundling reduces their effective access cost by more. Empirical Prediction II maps to the cross-category heterogeneity in Section 4.3.1: categories more commonly bundled with grocery trips, or with higher within-category substitution elasticities, exhibit larger treatment effects.

4 Empirical Results

We organize the results in three parts. First, we validate the matching procedure and the recorded opening dates (Sections 4.1–4.2). Second, we estimate demand spillovers to incumbent businesses as measured by foot traffic (Section 4.3), documenting heterogeneity by business category, grocery store type, and the pre- versus post-COVID divide. Third, we examine whether these demand spillovers translate into real business formation and employment growth (Section 4.4), decomposing the effects into entry, exit, and incumbent adjustment margins.

¹⁵Note the standalone term for Treat_{ij} alone is collinear with the fixed effect α_{ij} , hence omitted in the regression. The term for $D_{\tau,ijnt}$ alone is collinear with the local time trend ϕ_{inkt} and therefore omitted as well.

4.1 Validation of Predictive Model and Propensity Scores for Treatment

We begin by validating the estimated propensity score model used to identify counterfactual (CF) sites for grocery store openings. Figure 2 shows the distribution of estimated propensity scores for real openings and their matched CF sites. Panels (a) and (b) indicate substantial overlap between the two groups in both the pre-COVID (2018–2019) and post-COVID (2020–2022) cohorts, suggesting that the model performs consistently across periods with different market conditions. Panel (c), which pools all openings, displays a similarly high degree of overlap, confirming that treated and control sites receive comparable predicted probabilities of opening. Complementing this evidence, Table 2 reports the balance checks for demographic characteristics within five miles and counts of incumbent business establishments by category within 0.5 miles. Mean differences between treated and control sites are small relative to the underlying dispersion of the covariates, and no variables display systematic imbalances. While a few business-count measures exhibit modest differences, none approach conventional thresholds for statistical significance. The standardized mean differences (SMDs) provide further support: all 22 covariates have $|\text{SMD}| < 0.10$ —the conventional threshold for adequate balance in propensity-score matching (Austin, 2009)—and the vast majority fall below 0.05. Taken together, the substantial overlap in propensity score distributions and the well-balanced covariate profiles suggest that the matched CF sites constitute a credible comparison group for the real openings.

Figure B4 offers additional validation. Figure B4a reports the realization rate of matched CF sites by half-year batch—the percentage subsequently realized as actual store openings at selected distance thresholds. A substantial proportion of matched CFs, particularly in earlier cohorts, are eventually realized as actual store openings, indicating that our matching procedure successfully identifies locations with a high likelihood of future entry. Figure B4b provides the complementary coverage perspective, showing for each opening cohort the percentage of real store openings located within the given distance of a prior matched CF site. Collectively, these panels confirm the predictive validity and spatial representativeness of our matched CF pool.¹⁶

4.2 Validation of Openings using Foot Traffic Measure

Next, we validate the timing of grocery store openings in our sample and examine whether all the openings indeed happened. We check whether there is a sharp increase in foot traffic visiting the location of the opening before and after the month of each grocery store opening. We use an event study of the following specification.

$$Y_{it} = \alpha_i + \gamma_t + \sum_{\tau=-4}^{12} \mu_{\tau} D_{\tau,it} + \varepsilon_{it}. \quad (4)$$

Here, the outcome variable Y_{it} is the monthly number of visits at the SafeGraph POI matched to

¹⁶All realization and coverage statistics are based on a restricted sample of real openings for the top 15 non-dollar store grocery brands in the U.S., truncated at the end of 2022. Expanding the sample to include additional brands or more recent data would likely result in higher rates.

a grocery store opening i . α_i denotes the store-opening fixed effects, controlling for time-invariant characteristics specific to each grocery store location, and γ_t denotes the fixed effects of the calendar year by month. The variable $D_{\tau,it}$ denotes the set of relative event time dummies, equal to one if Y_{it} is observed in calendar month t is τ months relative to its opening. Our coefficient of interest, which quantifies the effects of the grocery store on foot traffic at the opening location, is denoted by μ_τ .

Figure 3 plots estimated effects, along with 95% confidence intervals. We observe that in the four months before the grocery store openings, there are no significant pre-trends in foot traffic at the opening location. Between 6 and 12 months after the opening, average foot traffic increases by approximately 691 visits relative to the pre-opening period. This represents a 404% increase relative to the pre-period median monthly visits of 137. The sharp rise in visits following openings, coupled with the absence of pre-trends, confirms that our identified openings correspond to the actual opening events and that the recorded opening dates are accurate.¹⁷

4.3 Treatment Effects on Foot Traffic to Surrounding Businesses

We now estimate the causal effect of grocery store openings on foot traffic to surrounding businesses using the event-study specification in Equation (3), with the log of monthly visits as the outcome. For the area surrounding each real opening and its matched CF sites, we use a distance band of 0.025 miles to define concentric rings in which surrounding businesses are located. Our regression sample consists of a balanced panel of all incumbent business establishments (SafeGraph POIs) within 0.1 miles¹⁸ of each real opening site and from each matched CF site, observed from 4 months before the openings to 12 months after the openings. Figure 4 plots the average of the estimated coefficients, $\beta_{\tau,n}$, for each concentric ring in the first 12 months after the openings. The figure thus summarizes average treatment effects in percent terms on foot traffic visits to nearby businesses after the grocery store openings. These spillover effects are highly localized as they decay to be statistically indistinguishable from zero for businesses located further than 0.1 miles from the actual and CF sites. Within 0.1 miles, the businesses surrounding the real openings experience a significant increase in foot traffic, relative to those surrounding the CF sites. Due to the localized nature of spillovers, we subsequently focus our analysis on comparing outcomes of businesses within 0.1 miles.¹⁹

We adapt equation (3) to define a single ring made of all businesses within 0.1 miles of either the real opening or its matched CF sites. Figure 5 plots the estimated coefficients β_τ for the effects

¹⁷The opening locations typically exhibit a few months of vacancy between tenancies, so foot traffic from the previous tenant does not confound the post-opening treatment effects on surrounding businesses.

¹⁸The geographic range of 0.1 miles we use to define our treatment groups is supported by the precision that SafeGraph can identify a visit to a POI. Details of SafeGraph algorithms that cluster GPS points to a single place can be found in their white paper: <https://www.safegraph.com/guides/visit-attribution-white-paper>.

¹⁹The spatial decay pattern documented in Figure 4 is consistent with the empirical predictions of the conceptual framework described in Section D. In the model, foot traffic to a nearby business is a function of the effective travel time required to visit that business. Following a grocery store opening, effective travel time falls more for businesses located closer to the new store due to trip bundling, leading to larger increases in foot traffic for nearby establishments.

on log monthly foot traffic for $\tau \in \{-4, 12\}$ months relative to openings. We confirm that there are no differential trends in foot traffic to nearby businesses in the treatment and control groups before the opening of the grocery stores. We find that the openings of grocery stores have economically and statistically significant positive effects on surrounding businesses. Average foot traffic in the treatment group increases by 23.4 percent 6–12 months after opening.²⁰

One concern for identifying the causal spillovers on foot traffic could be that some developers and shopping center owners consider these spillover effects when endogenously determining property tenant mix at the same time as introducing the grocer anchor (Jardim, 2016). In particular, Publix has been increasing its ownership of shopping centers where it operates to control the neighboring tenant mix.²¹ In this case, nearby foot traffic growth may not be due to the causal effect of the grocer anchors alone but rather because the co-tenants are endogenously determined. We perform a robustness check where we exclude the 85 openings of Publix from our sample. Figure B6 shows that the main results in Figure 5 remain robust.

4.3.1 Heterogeneous Effects on Foot Traffic by Category of Surrounding Businesses

We next examine how demand spillovers vary across business categories. We divide the surrounding businesses into six categories according to their NAICS codes: (1) Accommodation, Eating and Drinking (2) Finance, Real Estate, Communication, and Professional (3) Medical, Welfare, and Healthcare (4) Wholesale and Retail (excluding Grocery Stores) (5) Grocery Stores (6) Other Services.²²

To perform the heterogeneity analysis, we estimate the regression specification in equation (3) using the log of the monthly number of visits as an outcome and run the regression separately for each business category. Figure 6 summarizes the total effects of openings on foot traffic to nearby businesses of various categories. Nearby wholesale and retail stores (excluding grocery stores) experience an average increase of 30 percent in foot traffic during the 6–12 months following a grocery store opening. The strongest synergies are observed between grocery stores and wholesale and retail establishments, as well as hospitality services for accommodations, eating, and drinking. This is expected since these types of businesses most likely share a lot of common demand from the same group of customers who visit the shopping center where the grocery anchor is located. Clustering around a grocery store anchor thus helps attract a large pool of customers. By contrast, demand spillovers are smaller for other categories of businesses. For instance, medical, welfare, and healthcare services exhibit only a 10 percent increase in visits, corresponding to statistically significant treatment effects that are 20 percentage points lower than those for wholesale and retail stores. These businesses may not share the same customer base as grocery stores, as individuals

²⁰We also estimate equation (3) that traces the effects on log monthly foot traffic for event times τ up to 24 months after a grocery store opening. Figure B5 plots the estimated coefficients. The results indicate that treatment effects stabilize after 6–12 months and remain relatively constant through months 12–24.

²¹Publix owns nearly a third of its 1,167 stores, according to the company’s 2017 annual report.

²²We categorize businesses based on their NAICS codes. Table A2 provides detailed information on the construction of each business category.

typically visit their designated healthcare providers infrequently, limiting the scope for incidental foot traffic from grocery shoppers. As a result, synergies between grocery stores and healthcare services appear substantially weaker.

Interestingly, we find that grocery store openings in our sample, on average, do not appear to divert customers from nearby incumbent grocery stores. The estimated effect on foot traffic to other existing grocery stores within 0.1 miles is close to zero.²³

The heterogeneity across categories can be interpreted through the conceptual framework in Appendix D. In the model, heterogeneous treatment effects across business categories arise from two sources: (1) categories that are more commonly bundled with grocery trips experience larger reductions in effective travel time $\bar{t}_{kj}$, and (2) categories with a higher within-category elasticity of substitution σ_k translate a given travel-time reduction into a larger visit response. Because both channels operate at the category level, comparing the magnitudes of estimated coefficients across the category-specific regressions above is informative about the relative strength of trip bundling and substitution elasticities, even though we cannot separately identify the two without trip-level data.

4.3.2 Heterogeneous Effects on Foot Traffic by Type of Grocery Store Openings

We next examine the heterogeneous spillovers on surrounding businesses generated by the openings of different types of grocery stores. In Figure 7, we divide the grocery store openings into 3 types: national/regional grocery store chains (e.g., Whole Foods and Stop & Shop), big-box retailers (e.g., Walmart and Costco), and discounters (e.g., ALDI).²⁴ We find that national/regional grocery chains generate the strongest positive spillovers on surrounding businesses. On average, 6–12 months after the opening of a national/regional grocery store, foot traffic to surrounding businesses within 0.1 miles increases by 31 percent. Meanwhile, opening a big-box retailer or discounter increases foot traffic to nearby businesses by 16 percent and 17 percent, respectively.

To understand the sources of this heterogeneity, we examine three dimensions along which store types differ: neighborhood income, the composition of surrounding businesses, and timing of entry. Neither neighborhood income nor business composition systematically varies across store types in ways that would explain the observed spillover differentials. Specifically, although discounters are more likely to open in lower-income Census Block Groups (Table 3), estimated treatment effects show no systematic income gradient (Figure 8).²⁵ The composition of surrounding businesses is also similar across store types (Table 4). However, the timing of entry does differ: big-box retailers and discounters account for a disproportionately larger share of post-COVID openings (Table 5), with 508 of 851 openings occurring in 2020 or later versus 343 before COVID. We examine the

²³The difference in the effects between nearby incumbent grocery stores and wholesale and retail stores (excluding grocery stores) is economically large and statistically significant at about 30 percentage points.

²⁴Appendix Table A1 provides the detailed list of grocery store chains and the corresponding classification of store types.

²⁵Specifically, we rank the Census Block Groups that receive a grocery store opening by their median household income and assign each opening to one of four income quartiles. The ranking is therefore relative to the distribution of household income among opening locations, rather than to the national distribution of all Census Block Groups.

implications of this timing difference in the next subsection.

4.3.3 Heterogeneous Effects on Foot Traffic by Pre- and Post-COVID Openings

We further examine heterogeneity in treatment effects by splitting the sample into pre-COVID and post-COVID grocery store openings. Figure 9 reports estimates of equation (3) separately for the two subsamples. We find that pre-COVID openings generate substantially larger spillover effects: foot traffic at surrounding businesses increases by 42.7 percent 6–12 months after the openings. In contrast, for post-COVID openings, the estimated effects are considerably smaller, with foot traffic increasing by only 14.6 percent.

We consider two potential explanations for this attenuation in treatment effects. The first is a level channel: post-COVID grocery store openings may attract fewer in-person visits, which would directly limit the scope for spillovers to nearby businesses. A growing literature documents persistent changes in retail formats, consumer behavior, and location choices of grocery chains following the pandemic that are consistent with such a reduction in physical store traffic.²⁶

To assess the importance of this channel, we estimate equation (4) separately for the pre- and post-COVID subsamples. Figure 10 reports the corresponding estimates with 95% confidence intervals. We find that the increase in foot traffic to the grocery store opening locations themselves is substantially larger prior to COVID: the estimated effect during 6–12 months after opening is 1,090 visits in the pre-COVID sample, compared with 412 visits in the post-COVID sample, implying an effect that is approximately 165% larger before the pandemic. A simple back-of-the-envelope calculation suggests that this reduction in grocery store traffic alone could account for approximately 90% of the observed decline in spillover effects, assuming spillovers scale linearly with total foot traffic: the pre-COVID treatment effect on surrounding businesses is 2.92 times the post-COVID effect, while $(1 + 1.65)/2.92 = 90.6\%$.

A second, complementary possibility is that grocery trips are less often bundled with visits to nearby businesses post-COVID, so the same grocery traffic generates less incidental demand for adjacent retailers and restaurants. While we cannot test this mechanism directly—we do not observe consumer-level visit sequences—it is consistent with evidence that remote work relaxes time constraints, enabling consumers to decouple routine grocery runs from discretionary stops and travel farther for high-quality amenities (Qian and Su, 2024). In the conceptual framework of Appendix D, less trip bundling means a grocery opening delivers a smaller reduction in the effective cost of visiting nearby businesses, mechanically weakening spillovers even conditional on grocery traffic.

²⁶There are several potential explanations for the lower level of foot traffic, including changes in retail business models, shifts in the composition of store openings, and persistent changes in consumer behavior; see, for example, Monte et al. (2023) and Goolsbee and Syverson (2021), and the references cited therein.

4.4 Treatment Effects on Surrounding Business Dynamics

The foot-traffic results in Section 4.3 document demand spillovers to incumbent businesses. A natural next question is whether these demand spillovers translate into real reallocation—new establishments entering, marginal incumbents surviving, and local employment expanding. We examine both the extensive margin—entry and exit of establishments—and the intensive margin—expansion and contraction of continuing establishments in terms of employment.

For business i in period t , let E_{it} denote an outcome of interest, such as an indicator for active status, equal to one if the business is operational in period t and zero otherwise (before entry and after exit), or employment, which measures firm size. Following [Ellickson and Grieco \(2013\)](#), we define the growth rate as

$$g_{it} = \frac{E_{it} - E_{it-1}}{X_{it}}, \quad (5)$$

where $X_{it} = \frac{E_{it} + E_{it-1}}{2}$ is the average level of the outcome in periods t and $t-1$. The growth rate g_{it} therefore measures the change in the outcome from period $t-1$ to t . Under this definition, $g_{it} = 2$ for entry (i.e., $E_{it-1} = 0$ and $E_{it} > 0$) and $g_{it} = -2$ for exit (i.e., $E_{it-1} > 0$ and $E_{it} = 0$). We can then aggregate the growth rates at the individual business level to the group level according to the following formulas:

$$\begin{aligned} JCC_{kt} &= \frac{\sum_{i \in \mathcal{R}_k} X_{it} \max\{g_{it}, 0\} \mathbf{1}[g_{it} < 2]}{\sum_{i \in \mathcal{R}_k} X_{it}}, \\ JCE_{kt} &= \frac{2 \sum_{i \in \mathcal{R}_k} X_{it} \mathbf{1}[g_{it} = 2]}{\sum_{i \in \mathcal{R}_k} X_{it}}, \\ JDC_{kt} &= \frac{\sum_{i \in \mathcal{R}_k} X_{it} \max\{-g_{it}, 0\} \mathbf{1}[g_{it} > -2]}{\sum_{i \in \mathcal{R}_k} X_{it}}, \\ JDE_{kt} &= \frac{2 \sum_{i \in \mathcal{R}_k} X_{it} \mathbf{1}[g_{it} = -2]}{\sum_{i \in \mathcal{R}_k} X_{it}}. \end{aligned} \quad (6)$$

To illustrate, JCC_{kt} measures positive growth among continuing firms. It aggregates positive changes for firms with $g_{it} > 0$ that are not entrants ($g_{it} < 2$), weighting each firm by its average size X_{it} and normalizing by the total size of group k . The remaining components are defined analogously: JCE_{kt} captures growth due to entry, JDC_{kt} captures negative growth from contraction at continuing firms, and JDE_{kt} captures negative growth due to exit.

The overall growth rate of outcome E_{it} in group k is then

$$\begin{aligned} G_{kt} &= \frac{\sum_{i \in \mathcal{R}_k} X_{it} g_{it}}{\sum_{i \in \mathcal{R}_k} X_{it}} = \frac{\sum_{i \in \mathcal{R}_k} (E_{it} - E_{it-1})}{\sum_{i \in \mathcal{R}_k} X_{it}} \\ &= JCC_{kt} + JCE_{kt} - JDC_{kt} - JDE_{kt}. \end{aligned} \quad (7)$$

Thus, group-level growth decomposes into positive contributions from expansion of continuing firms and entry, and negative contributions from contraction and exit. When E_{it} denotes the

active-status indicator, G_{kt} corresponds to the growth rate in the number of businesses in group k . In this case, JCE_{kt} and JDE_{kt} can be interpreted as the shares of openings and closures relative to the stock of existing businesses.

In our setting, groups are defined by concentric distance rings around each opening site. Specifically, for each grocery opening i at site j , we partition the surrounding area into rings of width 0.1 miles, indexed by $n \in \{1, \dots, N\}$. Our regression sample consists of a panel of all businesses located within 0.5 miles of each real opening site and its matched counterfactual (CF) site, observed from three years before to three years after the opening.

To estimate the effects of grocery store openings on nearby business dynamics, we implement the following event-study specification:

$$Y_{ijnt} = \sum_{\tau \in [-s_1, s_2]} \beta_{\tau, n} (\text{Treat}_{ij} \times D_{\tau, ijt}) + \alpha_{inj} + \phi_{int} + \varepsilon_{ijnt}, \quad \forall n \in \{1, \dots, N\}. \quad (8)$$

The outcome Y_{ijnt} denotes the outcome of interest for ring n around site j associated with opening i in period t (e.g., the growth in the number of businesses within the ring). The indicator Treat_{ij} equals one if site j is a real opening site and zero if it is a matched CF site. The event-time indicator $D_{\tau, ijt}$ equals one if period t is τ periods relative to the opening of grocery i . The fixed effect α_{inj} controls for time-invariant characteristics specific to opening i , ring n , and site j , while ϕ_{int} denotes opening-by-ring-by-calendar-year fixed effects. The coefficients of interest, $\beta_{\tau, n}$, trace out the dynamic effects of an opening on nearby businesses in ring n at event time τ . We normalize $\beta_{-1, n} = 0$, so all estimates are relative to the period immediately preceding the opening. Standard errors are clustered at the opening-site level. This specification mirrors Equation (3) but uses annual frequency and wider distance rings (0.1 miles) to accommodate the lower frequency at which business formation and exit are observed in the Data Axle data.

Following the specification in Equation (8), Figure 11 presents the spatial pattern of treatment effects implied by the event-study estimates. Specifically, we average the post-opening coefficients over the first three years after entry and plot the resulting effects by distance ring. Panel (a) reports effects on the growth rate of the number of establishments, and Panel (b) reports effects on employment growth. The impacts are highly localized. Within 0.1 miles (ring 1), grocery entry increases establishment growth by 7.0 percentage points and employment growth by 8.8 percentage points relative to matched counterfactual sites. The effects attenuate sharply with distance. In the 0.1–0.2 mile ring (ring 2), the corresponding increases are 2.0 percentage points for establishments and 3.0 percentage points for employment. Beyond 0.2 miles, the estimates are small and statistically indistinguishable from zero. These results indicate that grocery anchors generate spatially concentrated spillovers that are largely confined to the immediate vicinity of the store, with little evidence of positive effects beyond 0.2 miles. The sharp spatial attenuation is consistent with a trip-chaining mechanism: consumers who visit the grocery store combine that trip with stops at nearby establishments, lowering the effective time cost of accessing those businesses (Pashigian and Gould, 1998; Miyauchi et al., 2025). Because trip-chaining is practical only within

walking distance, these demand-side spillovers decay rapidly beyond the 0.1–0.2 mile radius.

Figure 12 instead fixes distance at 0–0.1 miles, where effects are largest, and traces the full event-time path of the coefficients $\beta_{\tau,n}$ from Equation 8. Panel (a) shows that the establishment growth rate increases sharply after the opening, peaking one year later at 12.4 percentage points relative to control sites. Panel (c) documents a similar pattern for the employment growth rate, which peaks at 15.8 percentage points in the first post-opening year. In both outcomes, pre-trends are close to zero, and the growth rate differential returns to near zero by the second and third post-entry years. Because these coefficients capture the annual growth rate differential between treated and control sites, a return to zero indicates that the treated group’s growth rate converges back to the control group’s rate—not that the level of business activity declines. The cumulative gains in establishment counts and employment therefore continue to build over the full three-year post-entry window.

Together, the spatial and temporal profiles indicate highly localized effects on nearby business growth rates—concentrated within 0.1 miles and largest in the first post-entry year—that translate into persistent cumulative gains in establishment counts and employment. We next decompose the overall growth rate effects to identify the margins of adjustment.

Figure 13 further decomposes the overall treatment effects within 0–0.1 miles into contributions from business openings, closures, expansion of incumbents, and contraction of incumbents. The reported estimates correspond to the 0–3 year post-opening period relative to 1–3 years before entry. Panel (a) shows that the 7.0 percentage point increase in the establishment growth rate is almost entirely driven by entry and reduced exit. New business openings account for a 4.2 percentage point increase, while fewer closures contribute an additional 2.9 percentage points. The expansion and contraction components are mechanically zero in this panel, since the outcome is an indicator for active status and continuing establishments cannot “expand” or “contract” along this dimension. Panel (b) presents the analogous decomposition for employment growth, where all four margins can in principle contribute. The overall 8.8 percentage point increase in employment growth is primarily attributable to higher employment at entering firms (5.4 percentage points) and fewer business closures (2.9 percentage points). The contribution from incumbent expansion is modest (0.7 percentage points), and changes in contraction are economically negligible. Taken together, the decomposition indicates that the localized growth effects documented earlier are driven predominantly by establishment entry and reduced exit, rather than by intensive-margin expansion among incumbent firms. The modest contribution of incumbent employment expansion is not necessarily at odds with the substantial foot-traffic gains documented in Section 4.3: increased customer visits may raise revenue for incumbents without inducing them to hire additional workers, particularly if existing staff can absorb the higher volume. More broadly, the predominance of the extensive margin suggests that the demand created by the grocery anchor is captured primarily by new entrants who choose to locate nearby because of the demand opportunity, rather than through employment scaling at incumbents whose size may be constrained by pre-entry conditions.

4.4.1 Heterogeneous Effects on Business Dynamics by Business Category and Opening Type

Next, we examine whether the localized growth effects documented above vary across business categories and by the type of entering grocery store. Figure 14 displays the corresponding point estimates and confidence intervals, enabling us to assess how the magnitude of these spillovers depends on the sectoral composition of the neighborhood and the characteristics of the entrant.

For the analysis of heterogeneity by business category, we redefine group k in Equations (6) and (7) by interacting distance rings with the granular industry classifications listed in Table A2. We then estimate Equation (8) separately for each category within a 0–0.1 mile radius of the opening site. To identify the margins of adjustment, we decompose total growth into the respective contributions from firm entry and exit.

Panels (a) and (c) of Figure 14 present the estimated effects on establishment and employment growth. The most pronounced impacts are concentrated within the grocery sector itself: the establishment growth rate rises by 19 percentage points, driven by an 11-percentage-point increase in the entry rate and a 9-percentage-point reduction in the exit rate. The grocery employment growth rate is even more substantial at 21 percentage points. While the remaining five categories also exhibit positive responses, the magnitudes are more attenuated. Four of these sectors see growth rates ranging between 4 and 7 percentage points in both dimensions, whereas accommodation and food services show only marginal growth rate gains of 1 to 2 percentage points. Collectively, these patterns suggest that new grocery anchors generate significant localized spillovers, which are most salient among closely related retail segments. The comparatively muted response in accommodation and food services—despite the strong foot-traffic gains documented for this category in Section 4.3—likely reflects higher barriers to entry in the restaurant sector, including commercial kitchen buildout costs, health permits, and longer gestation periods that may extend beyond our three-year post-opening window.

Panels (b) and (d) of Figure 14 further explore heterogeneity by the type of the entering grocery store. We find that national and regional chains, as well as big-box and warehouse formats, catalyze the most significant localized effects on nearby business growth. Within the 0–0.1 mile radius, these types are associated with an approximate 9-percentage-point increase in establishment growth rates, while employment growth rates rise by 10 and 11 percentage points, respectively. In contrast, the effects of discount formats are substantially more muted, yielding impact magnitudes roughly half as large as those observed for their larger-format or chain counterparts.

Consistent with our decomposition results, these localized growth effects are driven primarily by increased entry and reduced exit, rather than expansion among incumbent establishments. This pattern aligns with theories of anchor tenants and retail agglomeration, whereby large-format or national chain stores serve as demand hubs that attract shared customer traffic and reduce effective search costs for nearby complementary businesses (Pashigian and Gould, 1998; Gould et al., 2005). Moreover, prior work emphasizes that store format shapes the strength of such spillovers, with larger-format retailers generating more pronounced consumer responses (Leung and Li, 2022). The

substantially weaker effects of discounters are consistent with several complementary channels: discount stores typically draw a more price-sensitive customer base less inclined toward discretionary companion shopping, operate smaller formats that generate fewer complementary shopping needs, and may carry weaker brand signals to prospective co-tenants about the local consumer profile. Taken together, our findings suggest that larger-format and chain grocery anchors generate stronger localized spillovers by concentrating foot traffic and fostering localized agglomeration economies.

4.4.2 Heterogeneous Effects on Business Dynamics by Industry

To further dissect the industry-level drivers of our baseline results, we estimate treatment effects disaggregated by 4-digit NAICS codes. Figure 15 presents point estimates for business growth and the corresponding contributions from entries and exits within 0–0.1 miles. For clarity, we focus on industries with significantly positive growth (at the 95% confidence level) that rank in the top two-thirds of the pre-existing industry share distribution. The industry-level heterogeneity reveals two distinct channels through which grocery anchors affect local business dynamics.

The first channel operates through direct demand spillovers. Panel (a) shows that the strongest extensive-margin responses in establishment counts occur in retail-adjacent and complementary consumption sectors—grocery stores, direct selling establishments, and restaurants—with total growth rates of approximately 7 percentage points, driven primarily by new entry. Panel (b) documents a parallel pattern for employment, with grocery stores and other information services showing the largest gains (8–9 percentage points). These are precisely the sectors that also experienced the strongest foot-traffic increases in Section 4.3, consistent with a trip-chaining mechanism.

The second channel is broader. Service industries such as depository credit intermediation and automotive repair also experience net entry, even though they did not exhibit comparably large measured demand spillovers. This pattern suggests that grocery anchor entry enhances local commercial viability more generally—perhaps by signaling neighborhood vitality or improving the local amenity environment—attracting businesses whose location decisions respond to the overall quality of the commercial cluster rather than to direct foot-traffic gains.

Taken together, the results in Section 4.4 complement and extend our foot-traffic analysis. While Section 4.3 demonstrated that grocery anchors generate localized demand spillovers to incumbent businesses, the findings here show that these demand effects translate into real reallocation: new establishments enter, marginal incumbents are sustained, and the local business ecosystem expands. The effects on growth rates are concentrated within 0.1 miles and largest in the first post-entry year, with cumulative gains in establishment counts and employment persisting through at least the third year after opening.²⁷

Mapping the results back to the decomposition in Section 3, the sharp spatial attenuation of both foot-traffic and business dynamics effects within 0.1 miles is consistent with the localized

²⁷Although we document COVID-related attenuation for foot-traffic spillovers in Section 4.3.3, the business dynamics analysis pools pre- and post-COVID openings because the Data Axle data are observed at annual frequency, leaving insufficient post-entry observations for cohort-specific event studies among the post-COVID subsample. We therefore interpret the business dynamics results as average effects across the full sample period.

spillover channel (TE^{local}) dominating the treatment-induced broad GE channel (GE^{treat}) at these distances. Under the conditions discussed in Section 3—namely, that area-wide shocks are absorbed by the matching design and that broad GE effects operate at spatial scales exceeding the distance between real and counterfactual sites—the estimated coefficients $\beta_{\tau,n}$ capture localized spillovers plus any differential in GE^{treat} between real and counterfactual neighborhoods.

5 Conclusion

This paper provides the first large-scale causal estimates of demand-side agglomeration spillovers from business openings in the non-tradable service sector. Using 851 grocery store openings across the United States (2019–2022) and a two-stage CNN–propensity-score identification strategy, we find that grocery openings increase foot traffic to nearby businesses within 0.1 miles by 23.4 percent, with spillovers decaying to zero beyond this radius. Effects are strongest for wholesale and retail stores (30 percent) and hospitality services, with national and regional chains generating the largest spillovers (31 percent). Demand spillovers attenuate sharply post-COVID—from 42.7 to 14.6 percent—consistent with remote work reducing multi-purpose shopping trips. Beyond incumbents, grocery anchors reshape the local business landscape: establishment growth rates rise by 7.0 percentage points and employment growth rates by 8.8 percentage points within 0.1 miles, driven by higher entry and lower exit rates.

These findings carry implications for commercial real estate investors, urban planners, and policymakers. The highly localized nature of spillovers suggests that tenant-mix decisions and zoning policies can materially affect local business formation and employment. The post-COVID attenuation further implies that remote work may durably reshape demand-side agglomeration forces in retail settings. Methodologically, our CNN–propensity-score pipeline demonstrates that deep-learning tools can be repurposed as spatial screening devices within a causal inference framework, an approach that may generalize to other spatial treatment-effect settings.

An important direction for future research is to examine whether these spillovers capitalize into commercial real estate outcomes such as rents and transaction prices. The commercial real estate transaction data currently available cover only a fraction of retail leases within the 0.1-mile radius around each site, leaving too few observations per site pair to detect capitalization effects with adequate statistical power in our matched-site framework. As data coverage of smaller retail leases expands, connecting demand-side agglomeration forces to asset pricing remains a promising avenue.

References

- Agrawal, Ajay and Iain Cockburn**, “The anchor tenant hypothesis: exploring the role of large, local, R&D-intensive firms in regional innovation systems,” *International Journal of Industrial Organization*, November 2003, *21* (9), 1227–1253.
- Athey, Susan and Guido W. Imbens**, “Machine Learning Methods That Economists Should Know About,” *Annual Review of Economics*, 2019, *11*, 685–725.
- Austin, Peter C.**, “Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples,” *Statistics in Medicine*, 2009, *28* (25), 3083–3107.
- , “Optimal caliper widths for propensity-score matching when estimating differences in means and differences in proportions in observational studies,” *Pharmaceutical Statistics*, 2011, *10* (2), 150–161.
- Basker, Emek**, “Job Creation or Destruction? Labor Market Effects of Wal-Mart Expansion,” *The Review of Economics and Statistics*, 2005, *87* (1), 174–183.
- Baum-Snow, Nathaniel, Nicolas Gendron-Carrier, and Ronni Pavan**, “Local Productivity Spillovers,” *American Economic Review*, 2024, *114* (4), 1030–1069.
- Benjamin, John D., Glenn W. Boyle, and C. F. Sirmans**, “Price discrimination in shopping center leases,” *Journal of Urban Economics*, November 1992, *32* (3), 299–317.
- Benmelech, Efraim, Nittai Bergman, Anna Milanez, and Vladimir Mukharlyamov**, “The Agglomeration of Bankruptcy,” *The Review of Financial Studies*, July 2019, *32* (7), 2541–2586.
- Bernstein, Shai, Emanuele Colonnelli, Xavier Giroud, and Benjamin Iverson**, “Bankruptcy Spillovers,” *Journal of Financial Economics*, September 2019, *133* (3), 608–633.
- Bertrand, Marianne and Francis Kramarz**, “Does Entry Regulation Hinder Job Creation? Evidence from the French Retail Industry*,” *The Quarterly Journal of Economics*, November 2002, *117* (4), 1369–1413.
- Billings, Stephen B., Shawn Rohlin, and Tingyu Zhou**, “The Death of Retail and Neighborhood Decline,” SSRN Scholarly Paper 4706738, SSRN 2024.
- Butts, Kyle**, “Difference-in-Differences Estimation with Spatial Spillovers,” *arXiv*, 2023. arXiv:2105.03737v3 [econ.EM], revised 10 Jun 2023.
- Clapp, John M., Stephen L. Ross, and Tingyu Zhou**, “Retail Agglomeration and Competition Externalities: Evidence from Openings and Closings of Multiline Department Stores in the U.S.,” *Journal of Business & Economic Statistics*, 2019, *37* (1), 81–96.
- Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker**, “Environmental Health Risks and Housing Values: Evidence from 1,600 Toxic Plant Openings and Closings,” *American Economic Review*, February 2015, *105* (2), 678–709.
- Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei**, “ImageNet: A large-scale hierarchical image database,” in “2009 IEEE Conference on Computer Vision and Pattern Recognition” IEEE 2009, pp. 248–255.

- Diamond, Rebecca and Tim McQuade**, “Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development,” *Journal of Political Economy*, June 2019, 127 (3), 1063–1117. Publisher: The University of Chicago Press.
- Dixit, Avinash K and Joseph E Stiglitz**, “Monopolistic competition and optimum product diversity,” *The American economic review*, 1977, 67 (3), 297–308.
- Duranton, Gilles and Diego Puga**, “Chapter 48 - Micro-Foundations of Urban Agglomeration Economies,” in J. Vernon Henderson and Jacques-François Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. 4 of *Cities and Geography*, Elsevier, January 2004, pp. 2063–2117.
- Ellickson, Paul B. and Paul L. E. Grieco**, “Wal-Mart and the geography of grocery retailing,” *Journal of Urban Economics*, May 2013, 75, 1–14.
- Ellison, Glenn and Edward L. Glaeser**, “Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach,” *Journal of Political Economy*, October 1997, 105 (5), 889–927.
- , – , and **William R. Kerr**, “What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns,” *American Economic Review*, June 2010, 100 (3), 1195–1213.
- Fung, Esther**, “Strip Centers Shine as Some Shoppers Sour on Malls,” *The Wall Street Journal*, January 2020. Published January 14, 2020.
- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther**, “Predictably Unequal? The Effects of Machine Learning on Credit Markets,” *The Journal of Finance*, 2022, 77 (1), 5–47.
- Gatzlaff, Dean, Stacy Sirmans, and Barry Diskin**, “The Effect of Anchor Tenant Loss on Shopping Center Rents,” *Journal of Real Estate Research*, January 1994, 9 (1), 99–110. Publisher: Routledge .eprint: <https://doi.org/10.1080/10835547.1994.12090738>.
- Glaeser, Edward L. and Joshua D. Gottlieb**, “The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States,” *Journal of Economic Literature*, December 2009, 47 (4), 983–1028.
- , **Hedi D. Kallal, José A. Scheinkman, and Andrei Shleifer**, “Growth in Cities,” *Journal of Political Economy*, 1992, 100 (6), 1126–1152.
- Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio**, “Generative adversarial networks,” *Communications of the ACM*, 2020, 63 (11), 139–144.
- Goolsbee, Austan and Chad Syverson**, “Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020,” *Journal of public economics*, 2021, 193, 104311.
- Gould, Eric D., B. Peter Pashigian, and Canice J. Prendergast**, “Contracts, Externalities, and Incentives in Shopping Malls,” *The Review of Economics and Statistics*, 2005, 87 (3), 411–422. Publisher: The MIT Press.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti**, “Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings,” *Journal of Political Economy*, 2010, 118 (3), 536–598.

- Henderson, J.Vernon**, “Marshall’s Scale Economies,” *Journal of Urban Economics*, January 2003, 53 (1), 1–28.
- Henderson, Vernon, Ari Kuncoro, and Matt Turner**, “Industrial Development in Cities,” *Journal of Political Economy*, October 1995, 103 (5), 1067–1090.
- Hu, Maggie Rong, Desmond Tsang, and Wayne Xinwei Wan**, “Corporate Relocation and Housing Market Spillovers,” *Management Science*, 2025, 71 (5), 4344–4376.
- Jardim, Eduardo**, “Essays on the Dynamic Decisions of Homeowners and Retail.” PhD dissertation, Duke University 2016.
- Jia, Panle**, “What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry,” *Econometrica*, 2008, 76 (6), 1263–1316.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan**, “Human Decisions and Machine Predictions,” *The Quarterly Journal of Economics*, 2018, 133 (1), 237–293.
- Knight, Samsun**, “Retail Demand Interdependence and Chain Store Closures,” SSRN Scholarly Paper 4234510, SSRN 2022.
- Koijen, Ralph S. J., Neel Shah, and Stijn Van Nieuwerburgh**, “The Commercial Real Estate Ecosystem,” Technical Report, CEPR Discussion Paper DP20053 2025. SSRN 5120847.
- Konishi, Hideo and Michael T. Sandfort**, “Anchor stores,” *Journal of Urban Economics*, May 2003, 53 (3), 413–435.
- Krishna, Priya**, “Why Do American Grocery Stores Still Have an Ethnic Aisle?,” The New York Times aug 2021. NYT link: <https://www.nytimes.com/2021/08/10/dining/ethnic-aisle-grocery-stores.html>.
- Krizhevsky, Alex and Geoffrey Hinton**, “Learning Multiple Layers of Features from Tiny Images,” Technical Report, University of Toronto 2009. Technical Report.
- Krugman, Paul**, “Increasing Returns and Economic Geography,” *Journal of Political Economy*, 1991, 99 (3), 483–499.
- Krugman, Paul R**, “Increasing returns, monopolistic competition, and international trade,” *Journal of international Economics*, 1979, 9 (4), 469–479.
- Kuiper, Niels, Mark van Duijn, and Arno van der Vlist**, “Retail Externalities and Distance in Shopping Malls,” in “European Real Estate Society (ERES) Conference Papers” 2021.
- Leung, Justin and Zhonglin Li**, “Big-Box Store Expansion and Consumer Welfare,” SSRN Scholarly Paper 3987727, SSRN 2022.
- Liu, Crocker and Peng Liu**, “Is What’s Bad for the Goose (Tenant), Bad for the Gander (Landlord)? A Retail Real Estate Perspective,” *Journal of Real Estate Research*, January 2013, 35 (3), 249–282. Publisher: Routledge eprint: <https://doi.org/10.1080/10835547.2013.12091369>.
- Liu, Crocker H., Stuart S. Rosenthal, and William C. Strange**, “The vertical city: Rent gradients, spatial structure, and agglomeration economies,” *Journal of Urban Economics*, July 2018, 106, 101–122.

- , —, and —, “Employment density and agglomeration economies in tall buildings,” *Regional Science and Urban Economics*, September 2020, *84*, 103555.
- Liu, Crocker, Yi Zheng, and Ling Zhu**, “Does Agglomeration Enhance Property Value?,” *Management Science*, 2026. Published online January 2026.
- Lorenz, Felix, Jonas Willkomm, Arne Handberg, and Michael Baguley**, “Interpretable Machine Learning for Real Estate Market Analysis,” *Real Estate Economics*, 2023, *51* (5), 1178–1208.
- Merriman, David, Joseph Persky, Julie Davis, and Ron Baiman**, “The Impact of an Urban WalMart Store on Area Businesses: The Chicago Case,” *Economic Development Quarterly*, November 2012, *26* (4), 321–333.
- Miyauchi, Yuhei, Kentaro Nakajima, and Stephen J Redding**, “The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data*,” *The Quarterly Journal of Economics*, 08 2025, *140* (4), 2507–2570.
- Monte, Ferdinando, Charly Porcher, and Esteban Rossi-Hansberg**, “Remote Work and City Structure,” Working Paper 31494, National Bureau of Economic Research jul 2023. Revised November 2025.
- Moretti, Enrico**, “Local Labor Markets,” in David Card and Orley Ashenfelter, eds., *Handbook of Labor Economics*, Vol. 4, Elsevier, 2011, pp. 1237–1313.
- Mullainathan, Sendhil and Jann Spiess**, “Machine Learning: An Applied Econometric Approach,” *Journal of Economic Perspectives*, 2017, *31* (2), 87–106.
- Neumark, David, Junfu Zhang, and Stephen Ciccarella**, “The Effects of Wal-Mart on Local Labor Markets,” *Journal of Urban Economics*, March 2008, *63* (2), 405–430.
- Oh, Ryungha and Jaeun Seo**, “What Causes Agglomeration of Services? Theory and Evidence from Seoul,” Technical Report oct 2023. Working paper.
- Pashigian, B. Peter and Eric D. Gould**, “Internalizing Externalities: The Pricing of Space in Shopping Malls,” *The Journal of Law and Economics*, April 1998, *41* (1), 115–142.
- Pollmann, Michael**, “Causal Inference for Spatial Treatments,” 2020. arXiv:2011.00373, revised Jan 2023.
- Qian, Franklin and Rose Tan**, “The Effects of High-skilled Firm Entry on Incumbent Residents,” SSRN Scholarly Paper 3982760, SSRN 2022.
- and **Yichen Su**, “Remote Work and Consumer Cities,” SSRN Scholarly Paper 4966864, SSRN 2024.
- Relihan, Lindsay**, “Is Online Retail Killing Coffee Shops? Estimating the Winners and Losers of Online Retail Using Customer Transaction Microdata,” CEP Discussion Paper 1836, Centre for Economic Performance, LSE 2022.
- Rosenbaum, Paul R. and Donald B. Rubin**, “Constructing a control group using multivariate matched sampling methods that incorporate the propensity score,” *The American Statistician*, 1985, *39* (1), 33–38.

- Rosenthal, Stuart S. and Joaquin A. Urrego**, “Eyes on the Street, Spatial Concentration of Retail Activity and Crime,” Technical Report dec 2023. Working paper (dated December 1, 2023).
- **and William C. Strange**, “Geography, Industrial Organization, and Agglomeration,” *The Review of Economics and Statistics*, 2003, *85* (2), 377–393.
- **and —**, “How Close Is Close? The Spatial Reach of Agglomeration Economies,” *Journal of Economic Perspectives*, August 2020, *34* (3), 27–49.
- Sadun, Raffaella**, “Does Planning Regulation Protect Independent Retailers?,” *The Review of Economics and Statistics*, December 2015, *97* (5), 983–1001.
- Shen, Lily and Stephen L. Ross**, “Information Value of Property Description: A Machine Learning Approach,” *Journal of Urban Economics*, 2021, *121*, 103299.
- Shoag, Daniel and Stan Veuger**, “Shops and the City: Evidence on Local Externalities and Local Government Policy from Big-Box Bankruptcies,” *The Review of Economics and Statistics*, July 2018, *100* (3), 440–453.
- Su, Yichen**, “Measuring the value of urban consumption amenities: A time-use approach,” *Journal of Urban Economics*, 2022, *132*, 103495.
- Wan, Wayne Xinwei and Thies Lindenthal**, “Testing Machine Learning Systems in Real Estate,” *Real Estate Economics*, 2023, *51* (3), 754–778.
- Waters, Shari**, “Choosing a Retail Store Location,” LiveAbout mar 2021. Updated on 03/12/21.
- Wheaton, William C.**, “Percentage Rent in Retail Leasing: The Alignment of Landlord–Tenant Interests,” *Real Estate Economics*, 2000, *28* (2), 185–204.
- Wolinsky, Asher**, “Retail Trade Concentration due to Consumers’ Imperfect Information,” *The Bell Journal of Economics*, 1983, *14* (1), 275–282.
- Zhou, Tingyu and John M. Clapp**, “The location of new anchor stores within metropolitan areas,” *Regional Science and Urban Economics*, January 2015, *50*, 87–107.






Figure 1: Locations of Grocery Store Openings


Notes: This figure plots the distribution of grocery store openings by CBSA in our sample. There are a total of 851 grocery store openings across the U.S. in our sample.



(a) Pre-COVID Sample



(b) Post-COVID Sample



(c) Full Sample

Figure 2: Distribution of Estimated Propensity Scores for Real Openings and Matched Counterfactual Sites

Notes: Panels (a) and (b) present the distribution of estimated propensity scores for actual openings and matched counterfactual locations, for the pre-COVID (2018-2019) and post-COVID (2020-2022) periods, respectively. Panel (c) shows the corresponding distribution for all periods combined. Each panel plots the estimated propensity score for the real openings and for their matched counterfactual sites.




Figure 3: Validation of Opening Dates using Foot Traffic Measure

Notes: Sample consists of all grocery stores opening in our sample. This figure plots the estimated coefficients of μ_τ of equation (4), along with 95% confidence intervals. μ_τ captures the effects of the grocery store opening on the foot traffic at the opening location. The outcome variable is the monthly number of visits. The standard errors are clustered at the grocery store level.




Figure 4: Treatment Effects on Foot Traffic to Surrounding Businesses as a Function of Distance

Notes: This figure plots the average treatment effects of grocery store openings on nearby businesses in each concentric ring in the first 11 months from the openings relative to the last 4 months before the openings by estimating equation (3). 95% confidence intervals are shown along with point estimates. The regression sample consists of a panel of all businesses within 0.2 miles of each of the grocery store opening sites and their matched CF sites, observed from 4 months before the openings to 12 months after the openings. For the area surrounding each real opening site and CF site, respectively, we use a distance band of 0.025 miles to define concentric rings in which surrounding businesses are located a distance of up to 0.1 miles from each site. Standard errors are clustered at the real or CF site level.




Figure 5: Treatment Effects on Foot Traffic to Surrounding Businesses Within 0.1 Miles

Notes: This figure plots the coefficients β_τ^g and corresponding 95% confidence intervals for each treatment group by estimating equation (3). Coefficients β_τ^g summarize the average effects of an opening on nearby businesses in treatment group g in period τ after the opening. The regression sample consists of a panel of all incumbent business establishments within 0.1 miles of the real grocery store openings and their matched counterfactual sites, observed from 4 months before the openings to 12 months after the openings. The treatment group consists of all businesses that are 0–0.1 miles from the real grocery store openings. The control group consists of all businesses that are 0–0.1 miles from the counterfactual sites for openings. Standard errors are clustered at the real or CF site level.




Figure 6: Heterogeneous Treatment Effects on Foot Traffic: by Category of Surrounding Businesses

Notes: This figure plots the heterogeneous treatment effects on the log of monthly visit counts to nearby businesses within 0–0.1 miles from the grocery store openings 6–12 months after opening relative to 1–4 months before opening. We report the treatment effects on different business categories by row. We categorize businesses based on their NAICS codes. Table A2 provides detailed information on the construction of each business category. We include “Grocery Stores” as a separate category. Examples of categories “Finance, Real Estate, Communication, and Professional” include real estate brokerages and banks. Examples of categories “Accommodations, Eating, and Drinking” include restaurants and bars. Examples of categories “Medical, Welfare, and Healthcare” include dentists and primary care providers. Examples of the category “Other services” include religious organizations and public administrations. Standard errors are clustered at the real or CF site level.




Figure 7: Heterogeneous Treatment Effects on Foot Traffic: by Grocery Store Type

Notes: This figure plots the heterogeneous treatment effects on the log of monthly visit counts to nearby businesses within 0–0.1 miles from the grocery store openings 6–12 months after opening relative to 1–4 months before opening. We report the treatment effects on different grocery store types by row. Big-box/club warehouses include Walmart, Costco, and Target; discounters include ALDI and Save-A-Lot; national/regional grocery chains include Publix, Whole Foods, and Kroger. See Appendix Table [A1](#) for the full classification. Standard errors are clustered at the real or CF site level.






Figure 8: Heterogeneous Treatment Effects on Foot Traffic: by Median Household Income of the Census Block Groups

Notes: We rank the Census Block Groups (CBGs) where the grocery store openings are located in terms of their median household income and divide the openings into quartiles. This figure plots the heterogeneous treatment effects on the log of monthly visit counts to nearby businesses within 0–0.1 miles from the grocery store openings 6–12 months after opening relative to 1–4 months before opening. We report the treatment effects by neighborhood income quartile by row. Standard errors are clustered at the real or CF site level.




(a) Pre-COVID Openings




(b) Post-COVID Openings

Figure 9: Heterogeneous Treatment Effects by Pre/Post COVID Openings

Notes: This figure plots the coefficients β_{τ}^g and corresponding 95% confidence intervals by estimating equation (3). Panel (a) plots the estimated results for the pre-COVID sample, and panel (b) plots the estimated results for the post-COVID sample. Pre-COVID openings are those occurring before 2020; post-COVID openings are those occurring in 2020 or later. The regression sample consists of a panel of all incumbent business establishments within 0.1 miles of the real grocery store openings and their matched counterfactual sites, observed from 4 months before the openings to 12 months after the openings. Standard errors are clustered at the real or CF site level.



(a) Pre-COVID Openings



(b) Post-COVID Openings

Figure 10: Foot Traffic before and after Opening: Pre-COVID and post-COVID Sample

Notes: This figure plots the estimated coefficients of μ_τ of equation (4), along with 95% confidence intervals for the pre- and post-COVID sample separately. μ_τ captures the effects of the grocery store opening on the foot traffic at that location. The outcome variable is the monthly number of visits. The standard errors are clustered at the grocery store level.






Figure 11: Treatment Effects on Surrounding Business Dynamics as a Function of Distance


Notes: This figure plots the average treatment effects of grocery store openings on nearby businesses in each concentric ring 0-3 years after opening relative to 1-3 years before opening by estimating equation (8). 95% confidence intervals are shown along with point estimates. The regression sample consists of a panel of all businesses within 0.5 miles from each of the grocery store opening sites and their matched CF sites, observed from 3 years before the openings to 3 years after the openings. For the area surrounding each real opening site and CF site, respectively, we use a distance band of 0.1 miles to define concentric rings in which surrounding businesses are located for a distance up to 0.5 miles from each site. Panel (a) shows the treatment effects on the growth of the number of businesses. Panel (b) shows the treatment effects on the growth of employment. Standard errors are clustered at the real or CF site level.




(a) Growth of Numbers of Businesses



(b) Growth of Numbers of Businesses: IPW




(c) Growth of Employment




(d) Growth of Employment: IPW

Figure 12: Treatment Effects on Surrounding Business Dynamics Within 0.1 Miles

Notes: This figure plots the coefficients $\beta_{\tau,1}$ and corresponding 95% confidence intervals by estimating equation (8). Coefficients $\beta_{\tau,1}$ summarize the average effects of an opening on nearby businesses within 0.1 miles in period τ after the opening. The regression sample consists of a panel of all businesses within 0.1 miles from the real grocery store openings and their matched counterfactual locations, observed from 3 years prior to the openings to 3 years after the openings. The treatment group consists of all businesses that are 0–0.1 miles from the real grocery store openings. The control group consists of all businesses that are 0–0.1 miles from the counterfactual locations for openings. Standard errors are clustered at the real or CF site level. Panel (a) & (c) plot the results of numbers of businesses and employment, respectively. Panel (b) & (d) additionally implement an inverse propensity score weighting.




(a) Growth of Numbers of Businesses




(b) Growth of Employment

Figure 13: Decomposition of Treatment Effects on Surrounding Business Dynamics


Notes: This figure presents the estimated treatment effects on surrounding business growth, decomposed into the respective contributions from openings, expansions, closures, and contractions. The analysis focuses on businesses within a 0–0.1 mile radius of the opening site, with the post-treatment window spanning 0–3 years after opening relative to a three-year pre-opening baseline. Panel (a) and (b) report the decomposition for the growth in the number of establishments and total employment, respectively. Point estimates are displayed with 95% confidence intervals. Standard errors are clustered at the actual or counterfactual (CF) opening site level.




(a) Growth of Numbers of Businesses:
by Surrounding Business Category



(b) Growth of Numbers of Businesses:
by Opening Grocery Store Type




(c) Growth of Employment:
by Surrounding Business Category




(d) Growth of Employment:
by Opening Grocery Store Type

Figure 14: Heterogeneity Treatment Effects on Business Dynamics by Business Category and Grocery Store Types

Notes: This figure reports heterogeneous treatment effects on surrounding business growth within 0–0.1 miles of grocery store openings over the 0–3 year post-opening period relative to 1–3 years prior to entry. Effects are shown separately by surrounding business category and by the type of grocery store opening, and are further decomposed into contributions from business entry (openings) and exit (closures). Panels (a) and (b) present results for the growth in the number of establishments, with business categories defined according to Table A2. Panels (c) and (d) present the corresponding results for employment growth. For heterogeneity by business category, group k in equations (6) and (7) is defined by interacting distance rings with business categories, and equation (8) is estimated separately for each category within 0.1 miles of the site. Point estimates are shown with 95% confidence intervals. Standard errors are clustered at the real or matched counterfactual site level.



(a) Growth of Numbers of Businesses



(b) Growth of Employment

Figure 15: Heterogeneous Effects on Business Dynamics by 4-Digit NAICS Industry

Notes: This figure reports heterogeneous treatment effects at the 4-digit NAICS industry level on surrounding business growth within 0–0.1 miles of grocery store openings during the year of entry and the first year thereafter, relative to 1–3 years prior to opening. Effects are decomposed into contributions from firm entry (openings) and exit (closures). Point estimates are shown with 95% confidence intervals. Industries are ordered from top to bottom according to the magnitude of the estimated total growth effect. The numbers in parentheses beneath each industry label report, respectively, the industry’s pre-opening share among surrounding businesses and its average post-opening growth rate at matched counterfactual sites. The figure displays industries with statistically significant positive effects at the 95% level and with pre-opening shares ranking in the top two-thirds of the distribution. Panel (a) presents effects on the growth in the number of establishments, and Panel (b) presents effects on employment growth. Standard errors are clustered at the real or matched counterfactual site level.

Variable	Category	Source
Population	Neighborhood Demographics	ACS
Share Under 16	Neighborhood Demographics	ACS
Share 16–65	Neighborhood Demographics	ACS
Share White	Neighborhood Demographics	ACS
Edu. College Plus	Neighborhood Demographics	ACS
Commute By Car	Accessibility/Traffic	ACS
Commute By Public Transit	Accessibility/Traffic	ACS
Commute Work From Home	Accessibility/Traffic	ACS
Commute Time 0-15	Accessibility/Traffic	ACS
Commute Time 15-30	Accessibility/Traffic	ACS
Commute Time 30-45	Accessibility/Traffic	ACS
Commute Time 45-60	Accessibility/Traffic	ACS
Commute Time 60-90	Accessibility/Traffic	ACS
Median HH Income	Customer Purchasing Power	ACS
Median Home Value	Customer Purchasing Power	ACS
Unemployment Rate	Customer Purchasing Power	ACS
Competitors	Competition Proximity	SafeGraph
Wholesale/Retail	Business Environment	SafeGraph
Finance	Business Environment	SafeGraph
Accom/Eat/Drink	Business Environment	SafeGraph
Medical Services	Business Environment	SafeGraph
Other Services	Business Environment	SafeGraph

Table 1: List of CNN Input Variables

Notes: This table lists all the input variables of the CNN model. We selected these variables to capture the primary factors that influence a grocery store’s location decision. ACS refers to the American Community Survey, five-year average statistics. For each half-year opening cohort, we use the corresponding five-year ACS estimates before the year of opening (e.g., the 2013–2017 ACS average for the openings in 2018, and the 2017–2021 ACS average for the openings in 2022). SafeGraph business point-of-interest (POI) data are accessed via the Dewey Platform. Variable abbreviations: Share Under 16 = share of population under age 16; Share 16–65 = share of population age 16–65; Edu. College Plus = share with above college degree; Commute By Car = share commuting by car; Commute By Public Transit = share commuting by public transportation; Commute Work From Home = share working from home; Commute Time XX = share with work commute time in each bracket (minutes); Median HH Income = median household income; Median Home Value = median home value. The “Competitors” category identifies grocery-related businesses; the remaining business environment categories are defined by broad NAICS sectors. Detailed NAICS code definitions for all business categories are provided in Appendix Table A2. The CNN model uses both the locations and counts of business POIs as inputs. Specifically, their latitude and longitude determine their positions in the input matrix, while their counts define the corresponding values in that matrix.

Variable	Treated Sites	Control Sites	Difference	SMD	T-statistics
<i>Demographic features within 5 miles</i>					
Population	3984.10 (8430.87)	4012.33 (6552.35)	-28.23	-0.00	-0.08
Share under 16	0.19 (0.06)	0.19 (0.06)	-0.00	-0.02	-0.48
Share 16–64	0.65 (0.08)	0.65 (0.08)	-0.00	-0.01	-0.15
Share White	0.72 (0.21)	0.72 (0.22)	0.00	0.01	0.25
Edu. College Plus	0.14 (0.10)	0.14 (0.11)	0.00	0.02	0.48
Commute By Car	0.84 (0.17)	0.84 (0.16)	0.00	0.01	0.27
Commute By Public Transit	0.05 (0.11)	0.05 (0.10)	-0.00	-0.00	-0.05
Commute Work From Home	0.06 (0.05)	0.06 (0.05)	-0.00	-0.02	-0.41
Commute Time 0–15	0.27 (0.15)	0.28 (0.16)	-0.01	-0.04	-0.83
Commute Time 15–30	0.36 (0.12)	0.36 (0.11)	-0.00	-0.01	-0.14
Commute Time 30–45	0.21 (0.09)	0.20 (0.10)	0.00	0.02	0.43
Commute Time 45–60	0.08 (0.05)	0.07 (0.06)	0.01	0.08	1.60
Commute Time 60–90	0.06 (0.06)	0.06 (0.05)	0.00	0.01	0.28
Median HH Income	70656.86 (29553.95)	69405.53 (32668.10)	1251.33	0.04	0.83
Median Home Value	275386.37 (198433.67)	279000.73 (230400.23)	-3614.36	-0.02	-0.35
Unemployment Rate	0.05 (0.04)	0.06 (0.04)	-0.00	-0.02	-0.42
<i>Business counts within 0.5 miles</i>					
Competitors	3.12 (5.96)	3.25 (6.17)	-0.14	-0.02	-0.46
Wholesale/Retail	23.63 (51.12)	21.97 (85.70)	1.66	0.02	0.49
Finance	39.84 (122.55)	36.50 (158.86)	3.34	0.02	0.49
Accom/Eat/Drink	32.92 (80.50)	25.74 (91.48)	7.18	0.08	1.72
Medical Services	35.13 (70.78)	29.66 (93.69)	5.47	0.07	1.36
Other Services	36.57 (93.00)	34.16 (120.70)	2.41	0.02	0.46

Table 2: Summary Statistics and Balance Checks

Notes: This table reports means and standard deviations (in parentheses) for treated sites (actual grocery-store openings) and matched CF sites as the control group. The treated sample consists of 851 opening sites, and the control group includes 2,089 unique CF sites. Each treated site is matched to *up to three* CF sites. To construct weighted balance statistics, CF sites are assigned inverse-match weights: when a treated site is matched to three CF sites, each CF receives a weight of 1/3; when matched to two CF sites, each CF receives a weight of 1/2, and analogously for other cases. Demographic variables are measured at the Census Block Group level using the latest ACS five-year estimates before the openings and are weighted by population overlap within a five-mile radius of each site, following the construction in equations (9) and (10). We measure business densities as the counts of SafeGraph POIs (opened at least one month before the focal opening) within a 0.5-mile radius, classified by NAICS sectors. The business category definitions match those used in Table 1, and we record them in Appendix Table A2. Standardized mean differences (SMDs) are computed as weighted standardized differences in means between treated and control groups, using the same inverse-match weights applied in the balance statistics. All t-statistics are from two-sided t-tests.

Grocery Type	Bottom 25%	25%–50%	50%–75%	Top 25%	Total
Big-box Retailers	21 (16.9%)	31 (25.0%)	31 (25.0%)	41 (33.1%)	124
Discounters	112 (28.5%)	109 (27.7%)	90 (22.9%)	82 (20.9%)	393
National/Regional Grocery Chains	63 (18.9%)	71 (21.3%)	100 (30.0%)	99 (29.7%)	334

Table 3: Distribution of Grocery Store Types by Median Household Income at the Census Block Group Level

Notes: This table shows the distribution of types of grocery store openings in each quartile from low-income to high-income. The neighborhood-level income is measured at the Census Block Group level.

Business Category	Big-box Retailer	Discounters	National/Regional Grocery Chains
Grocery	2.09%	3.18%	2.69%
Wholesale and Retail (excluding Grocery)	18.49%	20.90%	19.93%
FIRE	15.31%	16.63%	16.70%
Accommodations, Eating, & Drinking	17.26%	17.59%	17.65%
Medical, Welfare, and Healthcare	20.29%	15.84%	17.16%
Other Services	26.56%	25.87%	25.86%
Total	100.00%	100.00%	100.00%

Table 4: Composition of Surrounding Businesses by Grocery Store Type

Notes: This table reports the distribution of surrounding POIs by business category for each type of grocery store opening. Business categories include Grocery; Wholesale and Retail (excluding Grocery); Finance, Real Estate, Communication, and related services (FIRE); Accommodations, Eating, and Drinking; Medical, Welfare, and Healthcare; and Other Services.

Period	Big-box Retailer	Discounters	National/Regional Grocery Chains
Pre-COVID	47 (37.9%)	139 (35.4%)	157 (47.0%)
Post-COVID	77 (62.1%)	254 (64.6%)	177 (53.0%)
Total	124	393	334

Table 5: Timing of Grocery Store Openings Relative to the COVID Pandemic, by Store Type

Notes: This table reports the timing of the opening year relative to the COVID by grocery store opening. We classify grocery store openings by timing and define all stores opening in 2020 or later as post-COVID openings.

Appendix For Online Publication

A	Additional Tables	49
B	Additional Figures	51
C	CBG Demographic Weighting	58
D	Conceptual framework	59
D.1	Model setup	59
D.2	Travel time	60
D.3	Model Predictions	61

A Additional Tables

Grocery Store Chain	Openings
Big-Box Retailers	
Target	69
Costco	30
Walmart Supercenter	12
Other big-box retailers (with < 10 openings)	13
<i>Subtotal</i>	<i>124 (14.6%)</i>
Discounters	
ALDI	323
Grocery Outlet	39
Save-A-Lot	28
Other discounters (with < 10 openings)	3
<i>Subtotal</i>	<i>393 (46.2%)</i>
National/Regional Grocery Store Chains	
Publix Super Markets	85
Whole Foods Market	36
Sprouts Farmers Market	17
Harris Teeter	12
Safeway	12
Kroger	11
Others (with < 10 openings)	161
<i>Subtotal</i>	<i>334 (39.2%)</i>
Total	851 (100.0%)

Table A1: Grocery Store Classification and Openings by Type


Notes: This table lists the grocery store chains in our sample and the number of store openings for each chain during 2019–2022. Chains are classified into three types: big-box retailers, discounters, and national/regional grocery store chains. Chains with fewer than 10 openings are grouped into residual categories within each type. Subtotal rows report the number and percentage share of openings for each store type.

Business Category	NAICS Code Definition
Wholesale and Retail	42, 44, 45 (two-digit)
Finance, Real Estate, Communication, and Professional	52, 53, 54, 55, 56 (two-digit)
Accommodations, Eating, and Drinking	72 (two-digit)
Medical, Welfare, and Healthcare	62 (two-digit)
Other Services	51, 61, 71, 81, 92 (two-digit)
Grocery	445110, 445120, 452210, 452319 (six-digit)


Table A2: NAICS Code Definitions for Five Business Categories and Competitors

Notes: We follow [Miyachi et al. \(2025\)](#) in defining business categories of consumption amenities. For all business categories except Grocery, NAICS classifications are defined at the two-digit level, corresponding to broad industry sectors under the North American Industry Classification System (NAICS). For the Grocery category, we identify grocery-related businesses using six-digit NAICS codes: 445110 (Supermarkets and Other Grocery Stores), 445120 (Convenience Stores), 452210 (Department Stores, including firms such as Macy’s and Kohl’s), and 452319 (All Other General Merchandise Stores, including firms such as Walmart and Target).

B Additional Figures



(a) Data Augmentation Methods



(b) The three types of input samples for the CNN model

Figure B1: CNN Training Components: Data Augmentation and Input Types

Notes: This combined figure presents two core components of our CNN training setup—data augmentation and input sample types. Panel (a) demonstrates the three data augmentation methods—translation, rotation, and mirroring—used to enhance the diversity and robustness of the training samples. Translation shifts the spatial grid representation either vertically or horizontally, simulating alternative spatial placements. Rotation captures rotational invariance by reorienting the input matrices, reflecting the idea that once relative positions are fixed, absolute orientation should not affect locational context. Mirroring applies a left–right reflection, increasing the model’s exposure to spatially symmetric scenarios. Collectively, these augmentations improve the model’s generalization across diverse real-world settings. Panel (b) presents the three types of input samples for the CNN model. Type 1 includes at least one real opening, which we remove to construct an idealized counterfactual location. Type 2 includes a real opening. Type 3 is randomly selected on the map and may not contain any real openings. “Op” denotes a real opening. The 3×3 grid is shown for illustration; the CNN model uses 10×10 grids in practice.






Figure B2: Labeling Schemes for CNN Training: Traditional (One-Hot) vs. Heatmap

Notes: This figure contrasts the traditional one-hot labeling scheme with our heatmap-based labeling approach for the CNN model. In the traditional scheme, the label marks only the correct position and ultimately reduces the matrix to a single value. In the heatmap scheme, label values are assigned across the entire matrix to guide the CNN's learning, with the overall loss function optimized via heatmap regression. We assign a value of 1 to the exact location of a real opening and decreasing values to surrounding locations, with the decay determined by a Gaussian function $f(x, y) = e^{-A(x^2+y^2)}$. Here A denotes the amplitude, set to $A = 1$ in our model.




(a) Pair 1 (Ashburn, VA) — overall location comparison. Trader Joe's at One Loudoun (red dot) vs. matched CF at Sugarland Crossing (blue square); separation ~ 4.99 mi.




(b) Pair 1 — zoomed views. Left: One Loudoun with the Trader Joe's (red dot). Right: Sugarland Crossing CF parcel (blue square, 0.05 mi side), later realized as a Lidl (12 May 2021).

Figure B3 continued on next page.




(c) Pair 2 (Houston, TX) — overall location comparison. Target at Shepherd Square, Montrose (red dot) vs. matched CF at Stafford (blue square); separation ~ 12.36 mi.




(d) Pair 2 — zoomed views. Left: Target (opened 25 Oct 2020) within Shepherd Square, Montrose (red dot). Right: Stafford CF parcel (blue square, 0.05 mi side), later realized as a Costco Business Center (20 October 2021).

Figure B3: Two Real-Counterfactual Pairs: Relative Locations and Local Contexts

Notes: This figure presents two examples of real-counterfactual (CF) pairs—one from the pre-COVID batches and one from the post-COVID batches—illustrating how CF sites compare to the real openings they are matched with in our CNN model. Panels (a)–(b) show the Ashburn, VA pair: a Trader Joe’s opened on 6 Sep 18 at One Loudoun, while the matched CF site, located approximately 5 mi east at Sugarland Crossing, was later developed into a Lidl on 12 May 2021. Panels (c)–(d) show the Houston, TX pair: a small-format Target opened on 25 Oct 2020 at Shepherd Square in Montrose, while the matched CF site, about 12.36 mi southwest in Stafford, became a Costco Business Center on 12 Oct 2021. Overall maps (a) and (c) highlight inter-pair distances. The zoomed panels (b) and (d) depict the local surrounding contexts within a 0.5 mi radius, corresponding to the CNN’s primary feature-extraction range. In all maps, real openings appear as red dots and matched CF parcels as blue squares (0.05 mi per side), which matches the spatial resolution we use in the CNN.



(a) CF Realization Rate Over Time at Selected Distance Thresholds.



(b) Coverage of Real Openings by Prior CFs Across Distance Thresholds.

Figure B4: Validation of Matched Counterfactual Sites: Realization and Coverage Rates.

Notes: This figure assesses the validity of the matched counterfactual (CF) sites in capturing real store openings. Panel (a) evaluates the predictive validity of the matched CF sites by reporting, for the candidate CF sites matched to each half-year batch of real openings, the percentage subsequently realized as actual openings within selected distance thresholds (0.2 mi, 0.5 mi, 1 mi, and 2 mi). Panel (b) presents the converse perspective, documenting, for each half-year batch of real openings, the percentage that had a corresponding prior CF site within the same thresholds. Together, these measures assess both the quality of the matching approach and the representativeness of the CF pool in covering realized openings. All statistics are calculated within a restricted sample of real openings for the top 15 non-dollar store grocery brands in the U.S. and are truncated at the end of 2022. The inclusion of additional brands or more recent data may yield higher rates.




Figure B5: Treatment Effects on Foot Traffic to Surrounding Businesses Within 0.1 Miles

Notes: This figure plots the coefficients β_τ^g and corresponding 95% confidence intervals for each treatment group by estimating equation (3). Coefficients β_τ^g summarize the average effects of an opening on nearby businesses in treatment group g in period τ after the opening. The regression sample consists of a panel of all incumbent business establishments within 0.1 miles of the real grocery store openings and their matched counterfactual sites, observed from 4 months before the openings to 24 months after the openings. The treatment group consists of all businesses that are 0–0.1 miles from the real grocery store openings. The control group consists of all businesses that are 0–0.1 miles from the counterfactual sites for openings. Standard errors are clustered at the real or CF site level.




Figure B6: Robustness: Excluding Publix Openings

Notes: This figure shows the robustness of the main results after removing the grocery store openings of Publix. This figure corresponds to Figure 5 and plots the coefficients $\beta_{\tau,n}$ and corresponding 95% confidence intervals for each treatment group by estimating equation (3). Coefficients $\beta_{\tau,n}$ summarize the average effects of an opening on nearby businesses in the n -th ring in period τ after the opening. The regression sample consists of a panel of all incumbent business establishments within 0.1 miles of the real grocery store openings and their matched counterfactual sites, observed from 4 months before the openings to 12 months after the openings. The treatment group consists of all businesses that are 0–0.1 miles from the real grocery store openings. The control group consists of all businesses that are 0–0.1 miles from the counterfactual locations for openings. Standard errors are clustered at the real or CF site level.

C CBG Demographic Weighting

In the propensity-score matching stage, demographic features are computed as spatially weighted averages of Census Block Group (CBG) attributes. For each site s (either a real opening or a counterfactual) and each CBG c intersecting the five-mile radius of s , $B(s, 5)$, we define the weight

$$w_{c,s} = Pop_c \times \frac{\text{Area}(c \cap B(s, 5))}{\text{Area}(c)}, \quad (9)$$

where Pop_c denotes the population of CBG c . Then, for each demographic attribute f , the site-level covariate $D_{s,f}$ is given by

$$D_{s,f} = \frac{\sum_c w_{c,s} d_{c,f}}{\sum_c w_{c,s}}, \quad (10)$$

where $d_{c,f}$ is the value of attribute f in CBG c . This construction ensures that each CBG contributes to the site-level measure in proportion to its overlapping population.

D Conceptual framework

This section proposes a conceptual framework to guide our empirical exercises.

D.1 Model setup

There are K composite of service goods, and a numeraire good representing all non-service consumption. Each composite good is sold by a unique category of business. Following [Miyachi et al. \(2025\)](#), we categorize surrounding service businesses into 6 types: (1) wholesale and retail (excluding grocery); (2) grocery stores; (3) accommodations, eating, and drinking; (4) finance, real estate, communication, and professional services; (5) medical, welfare, and healthcare services; (6) Other services.

Each individual chooses the number of visits to each composite good, denoted as X_k , and a numeraire consumption good x_0 , and maximizes the following Cobb-Douglas utility function,

$$\max_{\{x_0, x_{kj}\}} x_0^{\theta_0} \prod_{k=1}^K X_k^{\theta_k}$$

where θ_k for $k = \{0, 1, \dots, K\}$ is the taste parameter for each type of composite good. We assume that the preference is homothetic so that $\sum_{k=0}^K \theta_k = 1$. Each composite good X_k is assumed to be a CES aggregator over an individual's visits to each place of interest (POI) j , denoted as x_{kj} .

$$X_k = \left[\sum_{j=1}^{J_k} x_{kj}^{\frac{\sigma_k-1}{\sigma_k}} \right]^{\frac{\sigma_k}{\sigma_k-1}}, \quad \sigma_k > 1$$

where J_k is the number of type- k POIs, and σ_k is the elasticity of substitution within business type k . We assume that $\sigma_k > 1$ for all k . This assumption implies that within the same category, each POI acts as a gross substitute for the others.

Each individual chooses the number of visits to type- k POI j , x_{kj} . Henceforth, we use the tuple (k, j) to denote type- k business j . The monetary cost of service offered by (k, j) is h_{kj} , which can be interpreted as the spending at (k, j) from a single visit. We assume that the spending per visit h_{kj} is constant before and after a grocery store entry.²⁸ In addition to this monetary cost of service, consumers incur a monetary cost corresponding to travel time, denoted as $\gamma \bar{t}_{kj}$, similar to [Su \(2022\)](#). Here, γ is the opportunity cost of the individual's time, and $\bar{t}_{kj}$ is the average time associated with visiting (k, j) , which we define more formally later. Therefore, the total price of purchasing service at POI j in category k is,

$$p_{kj} = h_{kj} + \gamma \bar{t}_{kj}$$

²⁸This assumption precludes the scenario where a POI upgrades its service quality to justify a higher price or charges a higher price markup over time.

We normalize the price of the numeraire good to be 1. Each individual is then subject to the following budget constraint

$$x_0 + \sum_{k=1}^K \sum_j^{J_k} x_{kj} (h_{kj} + \gamma \bar{t}_{kj}) \leq I,$$

where I is the individual's realizable income, if they were to have no visits to surrounding businesses and consume only the numeraire good x_0 .

Given the preference, the Dixit-Stiglitz price index for visits to type- k business is,

$$P_k = \left[\sum_j^{J_k} (h_{kj} + \gamma \bar{t}_{kj})^{1-\sigma_k} \right]^{1/(1-\sigma_k)} \quad (11)$$

The price index P_k reflects the unit cost of consuming the composite good X_k , aggregating service costs and travel costs across all POIs in category k (Dixit and Stiglitz, 1977). When $\sigma_k > 1$, access to more POIs lowers P_k , reflecting love-of-variety effects (Krugman, 1979). The magnitude of this variety effect is decreasing in σ_k : when POIs within a category are close substitutes, adding variety only weakly reduces the price index, but consumers are more responsive to price differences across individual POIs.

Consumer optimization indicates that the optimal number of visits to POI (k, j) can be expressed as,

$$x_{kj} = \left(\frac{p_{kj}}{P_k} \right)^{-\sigma_k} \times \frac{I\theta_k}{P_k} = (h_{kj} + \gamma \bar{t}_{kj})^{-\sigma_k} \times P_k^{\sigma_k-1} \times I\theta_k \quad (12)$$

From this equation, it is clear that individuals visit POIs with lower total prices more often. The elasticity of the number of visits to the price is governed by the elasticity of substitution σ_k .

D.2 Travel time

Recall that in the previous section, we denote the average travel time associated with visiting POI (k, j) as $\bar{t}_{kj}$. In this subsection, we formally define $\bar{t}_{kj}$, link it with individual trip bundling choice, and outline one assumption we make about the trip bundling.

Formally, let $\tilde{t}_{kj}$ denote the round-trip travel time from home to POI (k, j) and back, and let $\tilde{t}_{km \rightarrow k'n}$ denote the single-trip travel time between POI (k, m) and POI (k', n) . When an individual makes a single-stop trip to POI (k, j) , the associated travel time is simply $\tilde{t}_{kj}$. When an individual makes an N -stop trip bundle, denoted as trip i , following the route Home $\rightarrow (k_1, j_1) \rightarrow (k_2, j_2) \rightarrow \dots \rightarrow (k_N, j_N) \rightarrow$ Home, the total travel time for trip i is,

$$t_i = \frac{1}{2} \tilde{t}_{k_1 j_1} + \sum_{m=1}^{N-1} \tilde{t}_{k_m, j_m \rightarrow k_{m+1}, j_{m+1}} + \frac{1}{2} \tilde{t}_{k_N j_N}$$

For a bundled trip i on which POI (k, j) is one of the N stops, i.e. $(k, j) \in \{(k_1, j_1), \dots, (k_N, j_N)\}$,

we define the travel time allocated to (k, j) as,

$$t_{i,kj} = \tilde{t}_{kj} - \frac{1}{N} \left(\sum_{m=1}^N \tilde{t}_{k_m, j_m} - t_i \right),$$

where $\sum_{m=1}^N \tilde{t}_{k_m, j_m} - t_i \geq 0$ is the total time saved by bundling and $1/N$ allocates these savings equally across the N stops. For a single-stop trip to (k, j) , we have $t_{i,kj} = \tilde{t}_{kj}$. Let M_{kj} denote the total number of trips on which POI (k, j) is a stop. Taking the average across all M_{kj} trips, the average travel time associated with (k, j) is,

$$\bar{t}_{kj} \equiv \frac{1}{M_{kj}} \sum_{m=1}^{M_{kj}} t_{m,kj}$$

Feasibility condition. A trip bundle is feasible only if its total travel time does not exceed the sum of the standalone round trips to each stop:

$$t_i \leq \sum_{m=1}^N \tilde{t}_{k_m, j_m}.$$

This condition holds whenever POIs are spatially clustered so that successive legs of the chain are shorter than separate round trips from home—a consequence of the triangle inequality in geographic space.

Under the feasibility condition, the time savings from bundling are non-negative, so $t_{i,kj} \leq \tilde{t}_{kj}$ for every bundled trip. Since single-stop trips contribute $t_{i,kj} = \tilde{t}_{kj}$, averaging across all trips yields,

$$\bar{t}_{kj} \equiv \frac{1}{M_{kj}} \sum_{m=1}^{M_{kj}} t_{m,kj} \leq \tilde{t}_{kj}.$$

Moreover, the more frequently a POI is included in bundled trips, or the larger the time savings per bundled trip, the lower $\bar{t}_{kj}$ is relative to $\tilde{t}_{kj}$.

D.3 Model Predictions

A grocery store opening at location ℓ adds a new POI to category $k = \text{grocery}$, reducing the price index P_{grocery} through the variety channel. Crucially, it also creates new trip-bundling opportunities for nearby POIs in *all* categories: consumers visiting the new grocery store can combine that trip with stops at adjacent businesses, reducing the effective travel time $\bar{t}_{kj}$ for nearby POIs (k, j) . Both channels lower the effective cost of consuming local services.

From equation (12), we see that the log number of visits a POI (k, j) receives is,

$$\ln(x_{kj}) = -\sigma_k \ln(p_{kj}) + (\sigma_k - 1) \ln(P_k) + \ln(I\theta_k)$$

and consequently, comparing the log number of visits before and after the grocery store opening, we have,

$$\Delta \ln(x_{kj}) = -\sigma_k \Delta \ln(p_{kj}) + (\sigma_k - 1) \Delta \ln(P_k)$$

which indicates that changes in the number of visits to a POI are driven by two variables: (1) the change in its total price, and (2) the change in the overall price index for business category k .

Empirical Prediction I: In regressions where we control for opening case by business category by calendar year fixed effects, a POI closer to the new grocery store will experience a larger increase in the number of visits following the grocery store’s opening, compared to a POI in the same business category that is further away from the grocery store entry.

By conditioning on the opening case by business category and calendar year fixed effects, we directly control for both $(\sigma_k - 1) \Delta \ln P_k$ and $\ln(I\theta_k)$. Consequently, cross-sectional variation in $\Delta \ln(x_{kj})$ within a category is driven solely by variation in $\Delta \ln(p_{kj})$. Under our assumption that the service price h_{kj} is constant, the remaining variation in p_{kj} comes exclusively from changes in the effective travel time $\bar{t}_{kj}$. Trip bundling implies that POIs closer to the new grocery store experience a larger reduction in $\bar{t}_{kj}$, because they are more likely to be included in a multi-stop chain anchored by the grocery visit. Consequently, these nearby POIs experience a larger increase in the number of visits.

Empirical Prediction II: Comparing the average treatment effects *across* business categories (i.e., across separate category-specific regressions as in Section 4.3.1), categories whose services are more commonly bundled with grocery trips, or that have a larger within-category elasticity of substitution σ_k , will exhibit larger average increases in foot traffic.

The change in log visits for POI (k, j) after absorbing category-specific price-index movements is proportional to $\sigma_k \Delta \ln(p_{kj})$. Thus, the average treatment effect within category k depends on both (1) the typical reduction in effective travel time $\bar{t}_{kj}$ for POIs in that category (which is larger for categories more likely to be bundled with grocery trips) and (2) the elasticity of substitution σ_k (which amplifies the visit response to a given price change). Since σ_k is constant within a category, this prediction is testable by comparing the magnitudes of the estimated coefficients across category-specific regressions.