Supplementary Information for Urban retail contraction reveals a spatial vulnerability trap in essential retail access

Hossein Gazmeh[1] **and Xinwu Qian**[1,*]

[1]Department of Civil and Environmental Engineering, Rice University, Houston, TX
[*]xinwu.qian@rice.edu

Contents

1 Data sources, brand taxonomy, and variable definitions 3
2 Unfamiliarity and demographic comfort zones 6
3 Equilibrium existence, uniqueness, and solver convergence 7
4 Model estimation and training procedure 8
5 Welfare measure derivation 11
6 Fixed category budget assumption 11
7 Model validation and temporal generalization 12
8 Welfare loss across all 23 retail categories 14
9 Per-category and per-metro regression structure 15
10 Robustness of the spatial-versus-demographic decomposition 16
11 Channel decomposition of welfare loss 17
12 Vulnerability trap conditions across categories and metros 19
References 21

List of Tables

S1 Retail brand taxonomy. . . . 3
S2 ACS Variable definitions. . . . 5
S3 Metro-level demographic summary statistics. . . . 5
S4 Within-metro income quintile thresholds. . . . 5
S5 Illustrative demographic comfort zones for the grocery category. . . . 6
S6 Equilibrium solver convergence diagnostics. . . . 8
S7 Neural residual architecture. . . . 9
S8 Learned scalar parameters. . . . 10
S9 Parameter stability across training sample sizes. . . . 10
S10 Category budget share stability (January-November 2022). . . . 12
S11 Panel expansion diagnostics. . . . 12
S12 Temporal holdout performance across evaluation frameworks. . . . 13
S13 Per-metro temporal share performance. . . . 13
S14 Training versus held-out predictive performance. . . . 13
S15 Welfare loss by NAICS retail category under simultaneous 20% contraction of all 23 categories. . . . 14
S16 Per-category regression coefficients at 20% contraction. . . . 15
S17 Per-metro regression coefficients at 20% contraction. . . . 15
S18 Counterfactual gap closure under alternative income groupings. . . . 16
S19 Spatial counterfactual mechanism at 20% contraction. . . . 16
S20 Grocery channel decomposition by income decile at 20% contraction. . . . 17
S21 Channel decomposition by essential category and income quintile at 20% contraction. . . . 18
S22 $R_{\min}$ by income quintile across contraction levels. . . . 19
S23 Trap conditions under brand-preserving contraction. . . . 19
S24 Per-metro grocery gap monotonicity under traffic-based brand-preserving contraction. . . . 20

List of Figures

S1 Q1-Q5 welfare gap trajectories under brand-preserving contraction by essential category. . . . 20

1 Data sources, brand taxonomy, and variable definitions

Brand taxonomy. The dataset comprises 360 retail brands spanning 23 NAICS categories, with a total of 15,233 POIs across five California metropolitan areas. Brands were retained if they recorded at least 30 monthly visits and maintained consistent POI presence across all 36 months of the study period (January 2022 to December 2024). Of the 360 brands, 116 belong to the four essential retail categories used in the contraction simulations. Table S1 lists all 360 retail brands.

The four essential categories and their brand composition are as follows. Grocery Stores include 62 brands, led by 7-Eleven (1,071 POIs), Safeway (235), Smart & Final (180), Grocery Outlet (177), Ralphs (173), and Vons (172). Health and Personal Care Stores include 21 brands, led by CVS (687 POIs), Walgreens (440), Rite Aid (147), ULTA Beauty (147), Bath & Body Works (106), and Sally Beauty Supply (101). Gasoline Stations include 20 brands, led by Chevron (660 POIs), ARCO (296), Shell Oil (182), Speedway (125), Valero Energy (100), and Fastrip (66). General Merchandise Stores include 13 brands, led by Dollar Tree (373 POIs), Target (285), Dollar General (195), Costco (130), Walmart (104), and Family Dollar Stores (82). The remaining 19 non-essential NAICS categories (244 brands) include clothing, electronics, home improvement, restaurants, and specialty retail.

Table S1. Retail brand taxonomy. All 360 brands grouped by NAICS category, sorted by number of POI locations within each category. Numbers in parentheses indicate POI count in California. Essential categories used in contraction simulations are marked with ⋆.

Category (brands)	Brands (POIs)
⋆ Grocery Stores (62)	7-Eleven (1,071), Safeway (235), Smart & Final (180), Grocery Outlet (177), Ralphs (173), Vons (172), Stater Bros. Markets (160), Trader Joe's (139), Albertsons (123), Circle K (114), Sprouts Farmers Market (111), Whole Foods Market (88), Food 4 Less (77), Quik Stop (77), Raley's (58), Lucky Supermarkets (55), Save Mart Supermarkets (55), FoodMaxx Supermarkets (50), ALDI (49), Cardenas Markets (42), El Super (38), Vallarta Supermarkets (38), WinCo Foods (37), 99 Ranch Market (28), Superior Grocers (26), Pavilions (25), Bel-Air (20), Gelson's Markets (20), Foods Co. (17), Nob Hill Foods (15), Seafood City (15), H Mart (13), Island Pacific Seafood Market (13), Nugget Market (13), Holiday Market (10), Ray's Food Place (10), Barons Market (9), Bristol Farms (9), Lunardi's Markets (8), Mollie Stone's Market (8), Mother's Market (8), Super King Markets (8), Erewhon Natural Foods (7), Cost Less Foods (6), Lassens Natural Foods and Vitamins (6), Big Saver Foods (5), IGA (5), Marukai (5), Rancho San Miguel Markets (5), SavMor Foods (5), 168 Market (4), Arteagas Food Center (4), Mitsuwa Market Place (4), Numero Uno Markets (4), Super A Foods (4), Greatwall Supermarket (GW) (3), Jimbo's (3), Lions Market (3), Wild Fork (3), Chavez Supermarket (2), Supermercado Mi Tierra (2), ampm (2)
⋆ Health and Personal Care Stores (21)	CVS (687), Walgreens (440), Rite Aid (147), ULTA Beauty (147), Bath & Body Works (106), Sally Beauty Supply (101), GNC (General Nutrition Centers) (62), The Vitamin Shoppe (61), LensCrafters (59), Sephora (38), Sunglass Hut (19), Health Mart (16), Planet Beauty (13), Stanton Optical (13), LUSH (12), L'Occitane (7), Perfumania (7), Max Muscle (6), MAC Cosmetics (5), Le Labo (4), Nutrishop (4)
⋆ Gasoline Stations (20)	Chevron (660), ARCO (296), Shell Oil (182), Speedway (125), Valero Energy (100), Fastrip (66), 76 (39), Pilot Flying J (34), Rotten Robbie (31), Mobil (16), Exxon Mobil (15), Love's Travel Stops and Country Stores (13), Texaco (11), United Pacific (11), Marathon (9), Gulf Oil (8), Gas & Shop (3), Aisle 1 (2), Street Corner (1), TravelCenters of America (1)
⋆ General Merchandise Stores (13)	Dollar Tree (373), Target (285), Dollar General (195), Costco (130), Walmart (104), Family Dollar Stores (82), Five Below (76), DAISO (32), Sam's Club (29), Big Lots Stores (28), Navy Exchange (21), Curacao (7), FYE (4)
Automobile Dealers (16)	Toyota (29), Honda (25), CarMax (13), Chevrolet (10), Mercedes Benz (10), Audi (9), BMW (8), Kia Motors (8), Nissan (7), Pick n Pull (6), Subaru (6), Lexus (5), Volkswagen (4), Hyundai Motor (3), Mazda (2), Chevrolet,Cadillac (1)
Automotive Parts, Accessories, and Tire Stores (14)	O'Reilly Auto Parts (327), AutoZone (112), NAPA Auto Parts (86), America's Tire (63), Les Schwab (57), Pep Boys (48), Advance Auto Parts (47), Big O Tires (45), Firestone Complete Auto Care (29), Batteries Plus Bulbs (20), American Tire Depot (18), Discount Tire (10), Just Tires (8), Rent A Wheel (4)
Beer, Wine, and Liquor Stores (2)	BevMo (102), Total Wine & More (32)
Book Stores and News Dealers (3)	Barnes and Noble (41), Half Price Books (5), Deseret Book (2)
Building Material and Supplies Dealers (9)	The Home Depot (206), Ace Hardware (115), Harbor Freight Tools (108), Lowe's (106), Tractor Supply Co. (69), Sherwin-Williams (62), Dunn Edwards (35), DIY Home Center (6), AMES (4)

Category (brands)	Brands (POIs)
Clothing Stores (78)	Old Navy (75), Tilly's (75), Marshalls (70), dd's DISCOUNTS (68), Hot Topic (63), Claire's (58), LIDS (58), Zumiez (58), Burlington (56), The Men's Wearhouse (51), H&M (Hennes & Mauritz) (40), lululemon athletica (39), Victoria's Secret (38), Torrid (37), Banana Republic (34), The Children's Place (34), Nordstrom Rack (33), Forever 21 (30), Express (28), Gap (28), Windsor (28), PacSun (27), Anthropologie (24), Carter's (24), Chico's (21), Hollister (21), Levi Strauss & Co. (18), Urban Outfitters (18), White House Black Market (18), Charlotte Russe (17), Athleta (16), Cotton On (16), Michael Kors (16), Columbia Sportswear (15), Free People (14), Lane Bryant (14), Talbots (14), Calvin Klein (13), Destination XL (13), Francesca's (13), Pro Image Sports (13), Rainbow Shops (13), Ralph Lauren (13), Kate Spade (12), Nordstrom (12), Soma Intimates (12), Volcom (12), American Eagle Outfitters (11), J.Crew (11), J.Jill (11), Zara (11), Aéropostale (10), Buckle (10), Eddie Bauer (10), Janie and Jack (10), Madewell (10), maurices (10), Brandy Melville (9), Lucky Brand (9), Tory Burch (9), Work World (9), Abercrombie (8), Aritzia (8), Icing (8), Under Armour (8), Uniqlo (8), Jockey (7), Friar Tux Shop (6), Marine Layer (6), Nautica (6), Brooks Brothers (4), David's Bridal (4), Scrubs & Beyond (4), Gucci (3), LOFT (3), Louis Vuitton International (3), Ann Taylor (2), OshKosh B'gosh (2)
Department Stores (11)	Ross Stores (265), Kohl's (81), T.J. Maxx (73), Macy's (57), JCPenney (39), Saks Off Fifth (11), Bloomingdale's (10), Dillard's (3), Neiman Marcus (3), Saks Fifth Avenue (2), Sears (1)
Electronics and Appliance Stores (3)	Best Buy (96), Apple Retail Store (32), ZAGG (7)
Furniture Stores (7)	Ashley Furniture HomeStore (13), Bob's Discount Furniture (12), Living Spaces (10), Mattress Firm (9), Crate and Barrel (7), IKEA (7), CB2 (4)
Home Furnishings Stores (12)	HomeGoods (40), World Market (23), Floor and Decor (18), Williams-Sonoma (14), Pottery Barn (12), West Elm (12), At Home (10), Lamps Plus (10), The Container Store (8), Le Creuset (6), Restoration Hardware (6), Sur La Table (5)
Jewelry, Luggage, and Leather Goods Stores (7)	Daniel's Jewelers (25), Zales (21), Pandora (16), Don Roberto Jewelers (11), Tiffany & Co. (10), Coach (4), Fossil (1)
Lawn and Garden Equipment and Supplies Stores (1)	Armstrong Garden Centers (11)
Office Supplies, Stationery, and Gift Stores (10)	Staples (76), Party City (72), Hallmark Cards (43), BoxLunch (26), Spencer's (23), Office Depot (15), Paper Source (12), OfficeMax (7), Edible Arrangements (4), Disney Store (1)
Other Miscellaneous Store Retailers (9)	PetSmart (91), Petco (91), Pet Food Express (52), Leslie's Swimming Pool Supplies (38), Pet Supplies Plus (12), Unleashed by Petco (12), Centinela Feed & Pet Supplies (8), Kirkland's (6), Blick Art Materials (4)
Other Motor Vehicle Dealers (2)	Harley Davidson (15), West Marine (12)
Shoe Stores (22)	Vans (96), Boot Barn (48), Skechers (45), Foot Locker (35), Famous Footwear (34), DSW (Designer Shoe Warehouse) (31), WSS (31), Journeys (22), Kids Foot Locker (15), Nike (14), Red Wing Shoes (14), New Balance (12), ALDO (11), Crocs (11), Steve Madden (11), Adidas (9), Finish Line (9), Off Broadway Shoe Warehouse (8), Rack Room Shoes (8), ASICS (6), PUMA (6), UGG (3)
Specialty Food Stores (10)	See's Candies (101), Rocky Mountain Chocolate Factory (13), Ghirardelli Chocolate (10), It'sugar (8), Beach Hut Deli (7), Nespresso (6), Rocket Fizz (5), Duck Donuts (3), Lolli and Pops (3), Sugarfina (1)
Sporting Goods, Hobby, and Musical Instrument Stores (20)	Big 5 Sporting Goods (142), Michaels Stores (110), GameStop (99), Jo-Ann Fabric and Craft Stores (61), Hobby Lobby Stores (48), Dick's Sporting Goods (28), Cycle Gear (26), Build A Bear Workshop (22), Champs Sports (18), Sportsman's Warehouse (16), Guitar Center (15), REI (15), Road Runner Sports (15), Play It Again Sports (13), Lakeshore Learning Materials (12), Roger Dunn Golf Shops (10), The Golf Mart (6), Sports Fever (4), PGA TOUR Superstore (3), Golf Galaxy (2)
Used Merchandise Stores (8)	Goodwill Industries (70), The Salvation Army (26), Plato's Closet (11), Buffalo Exchange (5), Once Upon a Child (5), Uptown Cheapskate (4), Out of the Closet (3), Crossroads Trading (1)

Demographic and socioeconomic variables. CBG-level variables are drawn from the 2024 American Community Survey 5-Year Estimates[1]. Table S2 lists each variable, its ACS source table, unit, and role in the model. The primary variables entering the equilibrium model are the four racial composition proportions (White non-Hispanic, Black non-Hispanic, Hispanic, and Asian, from ACS table B03002) and log median household income (from B19013_001), which together define the demographic vectors used in the unfamiliarity term. Vehicle ownership (B08201) enters the regression analyses as a covariate capturing mobility constraints. Population density (B01003_001 / land area) is reported for descriptive comparison across metros but does not enter the equilibrium model or regression analyses.

Table S3 reports summary statistics for all demographic variables by metropolitan area. The five metros differ substantially in income (median household income ranges from $81,909 in Sacramento to $135,833 in San Jose), racial composition (percent Hispanic ranges from 21.4% in Sacramento to 44.1% in Los Angeles), and vehicle access (percent without vehicle ranges from 5.1% in San Diego to 9.6% in San Francisco). These differences motivate the within-metro z-scoring used in all pooled analyses.

Table S2. Variable definitions. Demographic and socioeconomic variables used in the model and regression analyses. ACS source refers to 2024 American Community Survey 5-Year Estimates table codes.

Variable	ACS Source	Unit	Transform	Role
Median HH income	B19013_001	$	log	Comfort zone, affordability
% White non-Hispanic	B03002	proportion	none	TV_{race} unfamiliarity
% Black non-Hispanic	B03002	proportion	none	TV_{race} unfamiliarity
% Hispanic	B03002	proportion	none	TV_{race} unfamiliarity
% Asian	B03002	proportion	none	TV_{race} unfamiliarity
% No vehicle	B08201	proportion	none	Regression covariate
Population density	B01003 / ALAND	per km^2	none	Descriptive
Brand count (15 km)	POI data	count	none	Spatial exposure
Mean dist. essential	POI data	km	log	Spatial exposure

Table S3. Metro-level demographic summary statistics. Values are mean ± SD across CBGs within each metropolitan area. The substantial differences across metros motivate within-metro z-scoring in all pooled analyses.

Variable	LA	SF	SD	SJ	SAC	Pooled
	n=5,899	*n=1,885*	*n=1,512*	*n=845*	*n=1,111*	*n=11,252*
Income ($K)	84±40	119±52	90±40	136±53	82±38	94±47
% White	30.1±26.3	37.8±24.4	47.2±26.0	31.7±20.4	53.1±24.5	36.1±26.6
% Black	6.2±11.7	7.5±10.9	4.3±6.7	2.3±3.8	6.5±8.7	5.9±10.4
% Hispanic	44.1±29.3	22.8±19.8	33.0±24.7	26.2±22.5	21.4±15.5	35.5±27.4
% Asian	16.1±17.8	25.8±21.2	10.6±13.1	35.3±22.8	12.7±14.0	18.1±19.2
% No vehicle	7.2±9.1	9.6±13.7	5.1±6.8	5.4±7.1	5.9±8.2	7.1±9.7
Pop. density (/km^2)	5,105±4,611	5,007±5,995	3,393±3,061	3,791±2,767	2,151±1,603	4,468±4,497

Within-metro income quintiles. Income quintiles are defined within each metropolitan area to ensure comparability across metros with different income distributions. Table S4 reports the upper income threshold for each quintile by metro. A CBG classified as Q1 (lowest income) in San Jose has median household income below $87,429, while Q1 in Sacramento requires income below $51,114. All regressions and welfare comparisons reported in the main text use within-metro quintile definitions.

Table S4. Within-metro income quintile thresholds. Upper bound of median household income ($) for each quintile, defined separately within each metropolitan area. Q5 has no upper bound.

Metro	Q1 upper	Q2 upper	Q3 upper	Q4 upper	Median
Los Angeles	$50,899	$67,285	$86,607	$111,983	$77,083
San Francisco	$73,969	$99,375	$125,142	$160,718	$111,847
San Diego	$56,128	$74,661	$94,536	$120,680	$84,831
San Jose	$87,429	$116,623	$146,762	$178,064	$131,310
Sacramento	$51,114	$66,923	$83,698	$109,279	$75,000

2 Unfamiliarity and demographic comfort zones

The unfamiliarity term $\Omega_{cb}(\sigma)$ in Equation 2 of the main text measures how demographically dissimilar a brand's visitor base is from a CBG's typical retail environment. This section details its construction.

Demographic comfort zone. For each CBG c and category k, the demographic comfort zone is the visit-weighted average demographic profile of all brands in category k that CBG c visited during the 2022-2023 training period. Formally, the comfort zone racial vector is

$$\text{comfort}^r_{ck} = \frac{\sum_{b \in \mathscr{B}_k} v_{cb} \cdot \tau^r_b}{\sum_{b \in \mathscr{B}_k} v_{cb}}, \quad r \in \{\text{White, Black, Hispanic, Asian}\},$$

where v_{cb} is the raw visit count from CBG c to brand b and τ^r_b is the observed proportion of brand b's visitors from racial group r. The comfort zone income is defined analogously using log median household income. All nonzero CBG-brand visit pairs are included with no minimum visit threshold. Comfort zones are computed once during preprocessing using all 17,731 California CBGs (not restricted to the five-metro study area) to maximize the demographic signal, and are stored as fixed matrices that do not update during equilibrium simulations, except when explicitly replaced in the demographic counterfactual (CF-B).

The four racial groups correspond to ACS table B03002 (Hispanic or Latino Origin by Race), using the categories White non-Hispanic, Black non-Hispanic, Hispanic, and Asian. These four groups account for approximately 96% of the California population. Because they represent proportions of a probability distribution, total variation distance is the natural metric for comparing racial composition vectors.

Income normalization. The income component of unfamiliarity is $D_{\text{income}} = \frac{2}{3}|\tau^{\text{inc}}_b(\sigma) - \text{comfort}^{\text{inc}}_{ck}|$, clamped to $[0,1]$. The scaling factor of $\frac{2}{3}$ (equivalently, dividing the absolute log-income difference by 1.5) is chosen so that 95% of CBG-brand income differences fall within $[0, 0.71]$, well below the clamp boundary. Across approximately 124,000 sampled CBG-brand pairs, the distribution of $|\tau^{\text{inc}}_b - \text{comfort}^{\text{inc}}_{ck}|$ has mean 0.438, standard deviation 0.329, and 95th percentile 1.069. Only extreme income mismatches (top 5%) approach the $[0,1]$ boundary. The TV_{race} component is inherently bounded to $[0,1]$ by the properties of total variation distance.

Equal weighting. The two components receive equal weight ($\frac{1}{2}$ each) in the composite unfamiliarity measure. This is an assumption-light choice given that both components are already scaled to $[0,1]$. The overall influence of unfamiliarity on utility is governed by the learned coefficient β_{unfam}, which enters the interpretable utility component as $+\beta_{\text{unfam}} \cdot \Omega_{cb}$ with $\beta_{\text{unfam}} = -0.47 < 0$ and is estimated end-to-end from the data.

Illustrative examples. Table S5 shows grocery-category comfort zones for five CBGs spanning the income distribution. A notable pattern emerges. Low-income CBGs (p10, log income = 10.69) have comfort zones with higher log income (11.21) than their own, reflecting that the grocery brands they visit serve a higher-income clientele on average. High-income CBGs (p90, log income = 11.98) have comfort zones with lower log income (11.35) than their own. The racial composition of comfort zones also varies across these examples, with lower-income CBGs showing higher Hispanic shares (0.468 at p10 versus 0.363 at p90) and lower White shares (0.347 versus 0.438), consistent with the residential segregation patterns in California metropolitan areas.

Table S5. Illustrative demographic comfort zones for the grocery category. Five CBGs spanning the income distribution. Comfort zone values are visit-weighted average demographic profiles of grocery brands each CBG historically visited. Low-income CBGs visit brands whose visitors have higher average income than the CBG itself, while high-income CBGs visit brands whose visitors have lower average income.

Pctl	CBG income	Comfort income	Δ	Cmf. White	Cmf. Black	Cmf. Hispanic	Cmf. Asian
p10	10.69	11.21	0.52	0.347	0.060	0.468	0.124
p25	11.00	11.25	0.25	0.385	0.059	0.419	0.138
p50	11.35	11.25	−0.10	0.521	0.055	0.297	0.126
p75	11.69	11.31	−0.38	0.417	0.053	0.387	0.143
p90	11.98	11.35	−0.63	0.438	0.051	0.363	0.148

Income values are in log dollars. Δ = comfort income − CBG income.

3 Equilibrium existence, uniqueness, and solver convergence

Existence and uniqueness of the QRE fixed point. We establish that the quantal response equilibrium[2] defined by Equation 3 of the main text admits a unique fixed point for fixed unfamiliarity under the conditions stated below. Consider the inner-loop operator T that maps an allocation vector σ to the softmax response given aggregate patronage $\mathbf{S}(\sigma)$. For fixed unfamiliarity Ω (as maintained within each inner-loop solve), the operator T maps the compact product of simplexes $\Delta = \{\sigma \geq 0 : \sum_{b \in \mathscr{B}_k} \sigma_{cb} = 1 \ \forall c, k\}$ into itself.

Proposition 1. *For fixed unfamiliarity Ω, finite utilities U_{cb}, positive choice precision $\lambda_k > 0$, nonnegative congestion sensitivity $\gamma_k \geq 0$, and a nonempty set of available brands $\mathscr{B}_k$, the inner-loop QRE admits a unique equilibrium allocation in each category.*

Proof. For each category k, consider the potential function

$$\Phi_k(\sigma) = \frac{1}{\lambda_k} \sum_{c \in \mathscr{C}_k^+} B_{ck} \sum_{b \in \mathscr{B}_k} \sigma_{cb} \log \sigma_{cb} - \sum_{c \in \mathscr{C}_k^+} B_{ck} \sum_{b \in \mathscr{B}_k} \sigma_{cb} U_{cb} + \gamma_k \sum_{b \in \mathscr{B}_k} \int_0^{S_b(\sigma)} \log(1+s)\, ds, \tag{S1}$$

where $\mathscr{C}_k^+ = \{c : B_{ck} > 0\}$ and $S_b(\sigma) = \sum_c B_{ck} \sigma_{cb}$. We adopt the convention $0 \log 0 = 0$, so the objective is continuous on the compact feasible set.

The entropy term is strictly convex in the allocation shares of CBGs with positive category budgets because $\lambda_k > 0$ and $B_{ck} > 0$. The utility term is linear, while the congestion term is convex because $\log(1+s)$ is increasing for $s \geq 0$ and S_b is linear in σ. Consequently, Φ_k is strictly convex on the feasible allocation set for positive-budget CBGs and admits a unique minimizer.

For any CBG c with $B_{ck} > 0$, the first-order condition with respect to σ_{cb}, subject to $\sum_{b \in \mathscr{B}_k} \sigma_{cb} = 1$, is

$$\frac{1}{\lambda_k}(1 + \log \sigma_{cb}) - U_{cb} + \gamma_k \log(1 + S_b) = \mu_c,$$

where μ_c is constant across brands for that CBG. The minimizer is interior for finite utilities, and rearranging the first-order condition yields

$$\sigma_{cb} = \frac{\exp(\lambda_k [U_{cb} - \gamma_k \log(1+S_b)])}{\sum_{b' \in \mathscr{B}_k} \exp(\lambda_k [U_{cb'} - \gamma_k \log(1+S_{b'})])},$$

which is exactly the QRE allocation in Equation 3. Hence the unique minimizer corresponds to the unique inner-loop fixed point for positive-budget CBGs. For CBGs with $B_{ck} = 0$, allocation shares do not affect aggregate patronage and are uniquely determined by the same softmax expression once the equilibrium patronage is known. □

This result establishes existence and uniqueness of the inner-loop equilibrium for fixed unfamiliarity. It does not establish that successive substitution is globally convergent, nor does it imply uniqueness of the full two-timescale system when unfamiliarity is endogenous.

Two-timescale solver. The system is solved with a nested iteration scheme. The inner loop solves the QRE (Equation 3) with unfamiliarity Ω treated as fixed, using pure successive substitution (no damping or momentum). The outer loop updates the visitor demographic profiles $\tau_b(\sigma)$ and recomputes unfamiliarity Ω_{cb} from the current allocation. The inner loop is run to convergence ($\|\sigma^{(t+1)} - \sigma^{(t)}\|_\infty < 10^{-6}$, maximum 100 iterations) before each outer-loop update. The outer loop terminates when $\|\sigma^{(t+1)} - \sigma^{(t)}\|_\infty < 10^{-4}$ or after 20 iterations.

The nested scheme is motivated by the two-timescale decomposition principle in stochastic approximation[3], applied here in a deterministic setting. Existence of an equilibrium for the coupled system can be established through the unfamiliarity update. For any fixed $\Omega \in [0,1]^N$, where N is the number of CBG-brand pairs, the inner-loop equilibrium $\sigma^*(\Omega)$ is unique and varies continuously with Ω because the strictly convex objective depends continuously on the fixed unfamiliarity values. Provided that each active category has a positive aggregate visit budget and the demographic comfort profiles are well defined, every available brand receives positive equilibrium patronage, so its visitor demographic profile is defined. The resulting update

$$F(\Omega) = \Omega(\sigma^*(\Omega))$$

is continuous and maps the compact convex set $[0,1]^N$ into itself. Brouwer's fixed-point theorem therefore guarantees at least one equilibrium of the full coupled system. This existence result does not establish uniqueness or convergence of the nested iteration. Convergence of the implemented solver is instead confirmed empirically across all training epochs and contraction simulations (Table S6).

Empirical convergence diagnostics. Across 147 training epochs and all contraction simulations, the solver never failed to converge. The inner loop typically converges in 5 iterations and the outer loop in 4 iterations. Table S6 reports convergence statistics from the training run (147 logged epochs). The iteration counts are consistent across epochs and across contraction scenarios (baseline, 10-50% contraction, recovery), indicating stable empirical convergence throughout the tested scenarios. The parameter summary also reports the worst-case value $\gamma_k \lambda_k / 4 \leq 0.133$ based on the learned parameter ranges. This quantity is descriptive and does not bound the full fixed-point Jacobian; it is therefore not used to establish contraction or solver convergence.

Table S6. Equilibrium solver convergence diagnostics. Statistics from 147 training epochs. The solver never failed to converge during training or any contraction simulation.

Parameter	Inner loop	Outer loop
Maximum iterations	100	20
Convergence tolerance (L^∞)	10^{-6}	10^{-4}
Typical iterations to convergence	5	4
Damping/momentum	None	None
Convergence failures	0 / 147 epochs	0 / 147 epochs
Learned category parameters		
γ_k range	$[0.489, 0.516]$	
λ_k range	$[0.972, 1.032]$	
Worst-case $\gamma_k \lambda_k / 4$ (descriptive)	0.133	

4 Model estimation and training procedure

Loss function. All model parameters are estimated jointly by penalized maximum likelihood on within-category visit shares aggregated over the 2022-2023 training period. For each CBG c and brand b in category k, the observed within-category share is $s_{cb}^{\text{obs}} = v_{cb}/B_{ck}$, where v_{cb} is the total visit count and $B_{ck} = \sum_{b' \in \mathscr{B}_k} v_{cb'}$ is the observed category budget. The model produces predicted shares $\hat{\sigma}_{cb}(\theta)$ through the equilibrium described in Equations 1-3 of the main text. The training objective is the categorical cross-entropy between observed and predicted within-category visit shares

$$\mathscr{L}(\theta) = -\frac{1}{|\mathscr{C}_{\text{valid}}|} \sum_{k=1}^{K} \sum_{c \in \mathscr{C}_k^+} \sum_{b \in \mathscr{B}_k} s_{cb}^{\text{obs}} \log \hat{\sigma}_{cb}(\theta) + \lambda_{\text{reg}} \|\alpha\|_2, \tag{S2}$$

where $\mathscr{C}_k^+ = \{c : B_{ck} > 0\}$ is the set of CBGs with positive budget in category k, $|\mathscr{C}_{\text{valid}}|$ is the total count of valid CBG-category pairs across all categories, α is the vector of brand fixed effects, and $\lambda_{\text{reg}} = 0.01$. CBG-brand pairs with zero observed visits contribute zero to the likelihood (since $0 \cdot \log \hat{\sigma} = 0$) and are not masked from the computation. CBGs with no visits to any brand in a category ($B_{ck} = 0$) are excluded entirely for that category. The loss is computed jointly across all 23 categories in each forward pass.

Forward pass. A single training step proceeds as follows. For each CBG-brand pair (c, b), five dyadic features are assembled from the preprocessed data: log distance to the nearest point of interest $\log d_{cb}$, income-affordability match $A_{cb} = 1/(1 + |\log y_c - \tau_b^{\text{inc}}|)$, log brand size $\log n_b$, log mean dwell time $\log \bar{\tau}_b^{\text{dwell}}$, and online transaction share o_b. The base interpretable utility $U_{cb}^{\text{interp,base}}$ is computed as the linear combination of these five features with their corresponding β coefficients plus a brand fixed effect α_b, with unfamiliarity added during equilibrium iterations (Equation 2). The neural residual $f_\theta(z_{cb})$ receives a 21-dimensional input z_{cb} that concatenates the same five scalar features with a 16-dimensional learned brand embedding e_b. The two components are combined as $(1-w) \cdot U_{cb}^{\text{interp,base}} + w \cdot f_\theta(z_{cb})$ through a learned mixing weight w (Equation 1), producing the base utility U_{cb}^{base}.

The base utility then enters the two-timescale equilibrium solver described in Supplementary Note 3. The outer loop computes unfamiliarity Ω_{cb} from the current allocation and updates effective utilities $U_{cb}^{\text{eff}} = U_{cb}^{\text{base}} + (1-w)\beta_{\text{unfam}} \cdot \Omega_{cb}$. The inner loop solves the QRE (Equation 3) with unfamiliarity treated as fixed, starting from uniform shares $\sigma_{cb}^{(0)} = 1/|\mathscr{B}_k|$ and iterating until convergence. Aggregate patronage $S_b = \sum_c B_{ck(b)} \sigma_{cb}$ is recomputed at every inner iteration to ensure that congestion feedback is captured endogenously. A typical training epoch requires approximately 4 outer iterations with 5 inner iterations each, for a total of approximately 20 fixed-point evaluations per epoch.

Gradient computation. The equilibrium is solved to convergence without gradient tracking. A single differentiable QRE step is then taken from the converged and detached equilibrium point σ^*, recomputing utilities from the current (differentiable) model parameters while holding the converged unfamiliarity and aggregate patronage fixed. This "differentiable

last iteration" approach[4,5] approximates implicit differentiation through the equilibrium fixed point while keeping training memory independent of the number of solver iterations. The approximation replaces the exact implicit derivative $d\sigma^*/d\theta = (I - \partial G/\partial\sigma|_{\sigma^*})^{-1}\,\partial G/\partial\theta|_{\sigma^*}$ with its first-order component $\partial G/\partial\theta|_{\sigma^*}$, where G is the QRE softmax operator. The unfamiliarity variable is also detached from the computation graph, so gradients flow through the QRE step but not through the demographic sorting feedback. This is a practical simplification given that endogenous unfamiliarity updates are small in magnitude (mean change of 0.008 at 20% contraction).

Neural network architecture. The neural residual f_θ is a feed-forward network with two hidden layers of 256 units each, ReLU activations, and a scalar output. The input is a 21-dimensional vector z_{cb} consisting of the five dyadic features concatenated with a 16-dimensional brand embedding e_b that is learned jointly with all other model parameters. The network uses Xavier uniform initialization for weights and zero initialization for biases. No dropout or normalization layers are applied. Brand embeddings are initialized as $e_b \sim \mathcal{N}(0, 0.01)$. Table S7 reports the layer-by-layer specification. The total model has 77,854 learnable parameters, of which 71,681 are neural network weights, 5,760 are brand embedding entries, and 413 are interpretable parameters (360 brand fixed effects, 6 β coefficients, 23 category-level λ_k, 23 category-level γ_k, and the mixing weight).

Table S7. Neural residual architecture. Layer-by-layer specification of the feed-forward network $f_\theta(z_{cb})$.

Layer	Type	Input dim	Output dim	Parameters
1	Linear + ReLU	21	256	5,632
2	Linear + ReLU	256	256	65,792
3	Linear	256	1	257
Neural subtotal				71,681
Brand embeddings (360×16)				5,760
Interpretable parameters				413
Total model parameters				**77,854**

Training procedure. The model is trained with the Adam optimizer[6] using initial learning rate 3×10^{-4}, weight decay 0.01, and default momentum parameters ($\beta_1 = 0.9$, $\beta_2 = 0.999$). Gradient norms are clipped at 5.0. The learning rate is halved when the training-set Spearman correlation plateaus for 15 consecutive epochs (ReduceLROnPlateau schedule). Training uses all 5,061 training CBGs simultaneously in each epoch (full-batch optimization) rather than mini-batching over subsets of CBGs. This design is required by the equilibrium structure, as the aggregate patronage $S_b = \sum_c B_{ck}\sigma_{cb}$ entering the congestion term requires all CBGs to be present for correct computation. All 24 months of visit data (January 2022 through December 2023) are summed into a single visit matrix before training. The model converges in 132 epochs (best checkpoint by training Spearman) out of a maximum of 147 epochs, with training loss decreasing smoothly from 2.131 to 2.118. Total training time is approximately 41 minutes on a single CPU. The near-identical performance between training and held-out CBGs ($\rho = 0.731$ versus 0.732 on visit counts, Table S14) suggests limited overfitting on this predictive metric, despite monitoring training-set performance.

All parameters are initialized from scratch ("cold start") without warm-starting from a simpler model. The six β coefficients are initialized at domain-informed values with the correct expected signs ($\beta_{\text{dist}} = -0.5$, $\beta_{\text{afford}} = +0.5$, $\beta_{\text{size}} = +0.5$, $\beta_{\text{dwell}} = +0.5$, $\beta_{\text{online}} = -0.5$, $\beta_{\text{unfam}} = -0.5$). Brand fixed effects are initialized using the log-share normalization $\alpha_b^{\text{init}} = \log(\hat{s}_b) - \overline{\log(\hat{s})}_{k(b)}$, where $\hat{s}_b$ is brand b's observed within-category market share and the mean is subtracted within each category. This initialization, standard in the BLP discrete choice literature[7,8], captures some of the cross-brand variation in visit shares before optimization begins. Category-level parameters are initialized at $\lambda_k = 0.5$ (via $\log\lambda_k = \log 0.5$) and $\gamma_k = 0.5$ (via $\log\gamma_k = \log 0.5$). The mixing weight is initialized at $w = 0.50$ (50% interpretable, 50% neural).

Two parameter constraints are imposed during training. Choice precision λ_k is parameterized as $\lambda_k = \exp(\log\lambda_k)$ with an upper clamp $\lambda_k \leq 1.10$ (enforced by clamping $\log\lambda_k \leq \log 1.10$ before exponentiation). This constraint limits the maximum choice precision during optimization. No lower bound is imposed on λ_k. Congestion sensitivity γ_k is parameterized as $\gamma_k = \exp(\log\gamma_k)$, ensuring positivity without an explicit clamp. No other parameters are constrained during optimization.

Learned parameters. Table S8 reports all scalar learned parameters with their initialized and final values. All six β coefficients retain their expected economic signs after unconstrained optimization. The mixing weight converges to $w = 0.37$, placing a weight of 0.63 on the interpretable component and 0.37 on the neural residual. Category-level parameters are tightly clustered (λ_k: mean 1.017, range [0.972, 1.032]; γ_k: mean 0.514, range [0.489, 0.516]), indicating that choice precision and congestion sensitivity are relatively homogeneous across the 23 retail categories. Brand fixed effects span a range of [−4.69,

+3.32] with mean approximately zero and standard deviation 1.43. A separate model trained on all 11,252 CBGs (used for the contraction simulations in Results 2-4) produces parameter estimates that agree with the 5,061-CBG model to three decimal places across all β coefficients, indicating stable coefficient estimates across the two training sample sizes (Table S9).

Counterfactual simulation procedure. After training, the full-CBG model (11,252 CBGs) is used for all contraction simulations with model parameters frozen. At each contraction level, closed POIs are identified by ranking all POIs within each essential category by observed visit traffic and removing the bottom $X\%$. For each brand b affected by closures, an accessibility scaling factor $a_{cb} = \sum_{j\in\text{open}}(d_{cj}+0.1)^{-1} / \sum_{j\in\text{all}}(d_{cj}+0.1)^{-1}$ is computed, where the sums run over open and all POIs of brand b respectively. This factor captures how much of brand b's spatial accessibility is retained from the perspective of CBG c after closures. The base utility is then adjusted as $U_{cb}^{\text{post}} = U_{cb}^{\text{pre}} + |\beta_{\text{dist}}| \cdot \log(a_{cb})$, which reduces utility in proportion to the lost accessibility with the same distance sensitivity learned from the data. If all POIs of a brand are closed, the brand is masked from the choice set entirely. The full two-timescale equilibrium is then re-solved from scratch using the same solver settings as training.

Table S8. Learned scalar parameters. All values from the production model trained on 5,061 CBGs with 2022-2023 visit data. The "Init" column reports the starting values used for cold-start optimization.

Parameter	Symbol	Init	Learned	Constraint
Distance decay	β_{dist}	-0.50	-0.97	None
Affordability match	β_{afford}	$+0.50$	$+0.47$	None
Brand size	β_{size}	$+0.50$	$+0.32$	None
Dwell time	β_{dwell}	$+0.50$	$+0.23$	None
Online share	β_{online}	-0.50	-0.07	None
Unfamiliarity	β_{unfam}	-0.50	-0.47	None
Mixing weight	w	0.50	0.37	$(0,1)$
λ_k (mean $\pm$ SD)		0.50	1.017 ± 0.022	≤ 1.10
γ_k (mean $\pm$ SD)		0.50	0.514 ± 0.006	> 0 (exp)
Brand FEs (mean $\pm$ SD)	α_b	log-share	0.00 ± 1.43	L_2 regularized

Table S9. Parameter stability across training sample sizes. Comparison of learned parameters from the 5,061-CBG spatial-holdout model and the 11,252-CBG full model used for contraction simulations. All β coefficients agree to three decimal places.

Parameter	5,061 CBGs	11,252 CBGs	$\lvert\Delta\rvert$
β_{dist}	-0.9714	-0.9712	0.0002
β_{afford}	$+0.4698$	$+0.4695$	0.0003
β_{size}	$+0.3240$	$+0.3242$	0.0002
β_{dwell}	$+0.2317$	$+0.2320$	0.0003
β_{online}	-0.0702	-0.0699	0.0003
β_{unfam}	-0.4694	-0.4691	0.0003
γ_k (mean)	0.5136	0.5137	0.0001
λ_k (mean)	1.0170	1.0172	0.0002
w	0.3704	0.3703	0.0001
Best epoch	132	133	1

Implementation and reproducibility. The model is implemented in Python 3.13 using PyTorch 2.9 (CPU build). Training and all simulations were performed on a single workstation equipped with an Intel Core i5-7600K (4 cores, 3.80 GHz) and 16 GB RAM, running Windows 10 Pro. A fixed random seed of 42 was applied to all random number generators (NumPy, PyTorch). The cold-start model (5,061 CBGs) trains in approximately 41 minutes and the full-CBG model (11,252 CBGs) in approximately 290 minutes. Key dependencies include NumPy 2.1, SciPy 1.15, and pandas 2.2. Source code is available at https://github.com/hogazme/DBED_Model_Sample.

5 Welfare measure derivation

Entropy decomposition. The share-weighted experienced welfare $W_{ck} = \sum_b \sigma_{cb} V_{cb}$ (Equation 4) and the inclusive value $\mathrm{IV}_{ck} = \frac{1}{\lambda_k} \log \sum_b \exp(\lambda_k V_{cb})$ are related through the Shannon entropy of the choice distribution. For a logit model with $\sigma_{cb} = \exp(\lambda_k V_{cb}) / \sum_{b'} \exp(\lambda_k V_{cb'})$, the inclusive value decomposes as

$$\mathrm{IV}_{ck} = W_{ck} + \frac{H(\sigma_{c\cdot})}{\lambda_k},$$

where $H(\sigma_{c\cdot}) = -\sum_b \sigma_{cb} \log \sigma_{cb}$ is the Shannon entropy of the within-category choice distribution. Taking the difference between post- and pre-contraction states and exponentiating yields

$$R_{ck} = R_{ck}^{\mathrm{IV}} \cdot \exp\left(-\frac{\Delta H_{ck}}{\lambda_k} \right),$$

where $R_{ck} = \exp(W_{ck}^{\mathrm{post}} - W_{ck}^{\mathrm{pre}})$ is the experienced welfare retention (Equation 5), $R_{ck}^{\mathrm{IV}} = \exp(\mathrm{IV}_{ck}^{\mathrm{post}} - \mathrm{IV}_{ck}^{\mathrm{pre}})$ is the inclusive-value retention, and $\Delta H_{ck} = H^{\mathrm{post}} - H^{\mathrm{pre}}$ is the change in entropy. This decomposition is a direct consequence of the logit structure and is implicit in the welfare formulas of McFadden[9] and Anderson, de Palma, and Thisse[10].

When brands close, changes in the distribution of visits can increase or decrease choice entropy. If entropy decreases ($\Delta H < 0$), the adjustment factor $\exp(-\Delta H / \lambda)$ exceeds 1, meaning experienced welfare retention exceeds inclusive-value retention. The inclusive value accounts for the entire available choice set, including brands that consumers rarely visit. Our share-weighted measure instead emphasizes payoffs at frequently visited brands, while the congestion adjustment captures the welfare reduction at surviving stores that absorb displaced demand.

Channel decomposition. To attribute welfare changes to access, congestion, and substitution, we hold demographic unfamiliarity at its pre-contraction value. Let $\widetilde{V}_{cb}^{\mathrm{post}} = V_{cb}^{\mathrm{pre}} + \Delta U_{cb}^{\mathrm{access}} - \gamma_k \Delta \log(1 + S_b)$ denote the post-closure payoff under this restriction, where $\Delta \log(1+S_b) = \log(1+S_b^{\mathrm{post}}) - \log(1+S_b^{\mathrm{pre}})$. Adding and subtracting $\sum_b \sigma_{cb}^{\mathrm{pre}} \widetilde{V}_{cb}^{\mathrm{post}}$ in Equation 4 gives

$$\begin{aligned} \Delta \widetilde{W}_{ck} = & \underbrace{\sum_{b \in \mathscr{B}_k} \sigma_{cb}^{\mathrm{pre}} \Delta U_{cb}^{\mathrm{access}}}_{\text{Access}} \\ & \underbrace{- \gamma_k \sum_{b \in \mathscr{B}_k} \sigma_{cb}^{\mathrm{pre}} \Delta \log(1 + S_b)}_{\text{Congestion}} \\ & \underbrace{+ \sum_{b \in \mathscr{B}_k} (\sigma_{cb}^{\mathrm{post}} - \sigma_{cb}^{\mathrm{pre}}) \widetilde{V}_{cb}^{\mathrm{post}}}_{\text{Substitution}}. \end{aligned} \tag{S3}$$

The three terms sum exactly to the change in experienced welfare evaluated with baseline unfamiliarity, using pre-contraction visit shares for access and congestion and post-contraction payoffs for substitution. The full counterfactual equilibrium still allows unfamiliarity to update endogenously.

6 Fixed category budget assumption

The model takes each CBG's monthly category budget B_{ck} as given from the data, following the standard two-stage budgeting framework in discrete choice demand models[11,12]. The Hicks composite commodity theorem provides conditions under which goods with proportionally changing prices can be aggregated[13]. Our model assumes fixed category visit budgets and estimates only the lower-level brand-choice stage.

This assumption affects the magnitude of the welfare findings. Holding the grocery budget fixed means that when a grocery store closes, consumers must find another grocery option within the category. In reality, some budget might shift to restaurants, convenience stores, or online delivery, meaning the actual within-category welfare loss could be smaller than what the model predicts. Fixed category budgets may therefore overestimate welfare losses when consumers can substitute effectively outside the modeled category.

The model's predictions depend primarily on the cross-sectional pattern of budgets (which CBGs allocate more to grocery versus gasoline) rather than on month-to-month fluctuations. This cross-sectional pattern is stable during the January-November 2022 observation period reported below. The mean grocery share across CBGs is 22.6-23.6% during January-November 2022,

with a cross-CBG coefficient of variation of 0.30-0.33 (Supplementary Table S10). Similar stability holds for health and personal care (HPC; 27.0-28.2%), gasoline stations (27.9-30.0%), and general merchandise (20.6-21.7%).

Table S10. Category budget share stability (January-November 2022). Mean category share of essential retail visits and cross-CBG coefficient of variation for each essential category. The stable cross-sectional pattern supports the fixed-budget assumption.

Category	Mean share (%)	Cross-CBG CV
Grocery	22.6-23.6	0.30-0.33
Health and Personal Care	27.0-28.2	0.41-0.44
Gasoline	27.9-30.0	0.33-0.36
General Merchandise	20.6-21.7	0.29-0.32

7 Model validation and temporal generalization

The main text reports spatial holdout performance and three counterfactual validity tests (Rite Aid closure comparison, local-competition test, and large-share brand removal). This section documents an additional dimension of model credibility, namely temporal generalization to 2024 data.

Panel expansion in the 2024 holdout data. Recorded GPS visit volume increased substantially between the training period (2022-2023) and the holdout period (2024), consistent with changes in panel coverage. Mean monthly visit volume increased from 11.3 million to 29.0 million (a 2.57× increase), while the number of observed CBG-brand pairs with positive visits decreased from 2.18 million (24 months aggregated) to 935 thousand (12 months), although these totals are aggregated over different observation windows and are not directly comparable as measures of panel coverage. The recorded visit growth is non-uniform across the income distribution and across metropolitan areas (Table S11). Higher-income neighborhoods experienced faster recorded visit growth (3.02× for Q5 versus 2.13× for Q1), and San Francisco grew 2.5 times faster than Los Angeles. Raw visit count comparisons between training and holdout periods are therefore unreliable, and we evaluate temporal holdout performance using within-category visit shares, which are invariant to uniform expansion within each CBG and category.

Table S11. Panel expansion diagnostics. Growth ratio of mean monthly visits per CBG between training (2022-2023) and holdout (2024) periods, by within-metro income quintile and by metropolitan area.

Panel A: By income quintile

Quintile	Train (monthly/CBG)	Holdout (monthly/CBG)	Growth ratio
Q1 (lowest)	925	1,967	2.13×
Q2	926	2,091	2.26×
Q3	987	2,538	2.57×
Q4	1,035	2,868	2.77×
Q5 (highest)	1,155	3,490	3.02×

Panel B: By metropolitan area

Metro	*N* CBGs	Growth ratio
Los Angeles	5,899	1.72×
San Francisco	1,885	4.32×
San Diego	1,512	3.28×
San Jose	845	2.70×
Sacramento	1,111	2.76×

Temporal holdout performance. Table S12 reports predictive performance under three evaluation frameworks. Within-category shares (Framework A) divide each CBG's visits to a brand by total visits to that brand's category, removing absolute-scale differences between periods. Budget-normalized counts (Framework B) scale 2024 visit counts to the training-period category budget for each CBG, removing CBG-by-category-level inflation while preserving cross-category variation. Raw counts (Framework C) compare model predictions directly to unadjusted 2024 visit counts.

Table S12. Temporal holdout performance across evaluation frameworks. Spearman rank correlations between predicted and observed CBG-brand visit patterns. Spatial holdout evaluates predictions on 6,191 held-out CBGs during the training period (2022-2023). Temporal holdout evaluates predictions on all 11,252 CBGs against 2024 observed data. Framework A uses within-category visit shares. Framework B uses category-budget-normalized counts. Framework C uses raw counts.

Model	Spatial (shares)	Temporal A (shares)	Temporal B (norm.)	Temporal C (raw)
Uniform	0.138	0.490	0.207	0.287
Gravity ($\alpha = 1.25$)	0.411	0.313	0.193	0.326
DBED	**0.600**	**0.594**	**0.409**	**0.473**

DBED achieves Spearman $\rho = 0.594$ on 2024 within-category shares, compared with 0.600 on the spatial holdout during the training period, and outperforms both baselines across all three temporal evaluation frameworks. Because the spatial and temporal evaluations cover different CBG samples, this comparison does not isolate a within-CBG change in predictive performance. A separate within-period temporal holdout (training on 2022, evaluating on 2023 within the same data source) yielded Spearman 0.624. Differences in temporal performance may reflect changes in retail availability, consumer preferences, and the composition of the mobility panel, which cannot be fully separated with the available data.

Per-metro temporal performance. Table S13 reports share-based spatial and temporal performance by metropolitan area. In the 2024 temporal holdout, correlations range from 0.571 in Los Angeles to 0.621 in San Diego, with comparable performance across all five metros.

Table S13. Per-metro temporal share performance. Within-category share Spearman correlations for DBED by metropolitan area.

Metro	*N* CBGs	Spatial	Temporal
Los Angeles	5,899	0.763	0.571
San Francisco	1,885	0.767	0.617
San Diego	1,512	0.760	0.621
San Jose	845	0.723	0.606
Sacramento	1,111	0.724	0.605
Pooled	11,252	0.756	0.594

Overfitting check. Table S14 compares performance on the 5,061 training CBGs and 6,191 held-out CBGs during the training period. DBED shows little evidence of overfitting on these predictive metrics. Held-out performance is marginally better than training performance on visit counts (Spearman 0.732 versus 0.731, $R^2 = 0.620$ versus 0.611), and essentially identical on within-category shares (0.600 versus 0.601).

Table S14. Training versus held-out predictive performance. All metrics computed on 2022-2023 visit data. Training and held-out correlations are nearly identical, suggesting limited overfitting on these metrics.

Model	Metric	Train (5,061)	Held-out (6,191)	Difference
DBED	Spearman (shares)	0.601	0.600	−0.001
	Spearman (visits)	0.731	0.732	+0.001
Gravity	Spearman (shares)	0.414	0.411	−0.003
	Spearman (visits)	0.439	0.436	−0.003
Uniform	Spearman (shares)	0.139	0.138	−0.001
	Spearman (visits)	0.321	0.322	+0.001

8 Welfare loss across all 23 retail categories

When category-proportional contraction is applied to all 23 NAICS retail categories simultaneously, removing the lowest-traffic 20% of POIs within each category and re-solving the full two-timescale equilibrium, every category experiences welfare loss (mean R_k ranges from 0.799 to 0.990; Table S15). Income-graded welfare loss is not unique to essential retail. Of the 23 categories, seven show significant negative β_{income} (low-income neighborhoods lose more, after controlling for spatial and demographic covariates), nine show significant positive β_{income} (high-income neighborhoods lose more), and seven are non-significant. Among the seven categories where low-income neighborhoods are disadvantaged, Grocery Stores exhibit by far the largest controlled income coefficient ($\beta = -0.0119$, $p < 0.001$), followed by Other Miscellaneous Store Retailers ($\beta = -0.0062$, $p < 0.001$) and Sporting Goods ($\beta = -0.0054$, $p < 0.001$). Uncontrolled Q1-Q5 gaps range from -5.36 pp (Sporting Goods) to $+3.63$ pp (Lawn and Garden Equipment), reflecting the diversity of spatial and demographic structures across retail sectors.

The essential-category results are effectively unchanged from the essential-only contraction reported in the main text (grocery gap -4.50 versus -4.51 pp, general merchandise -1.20 versus -1.20 pp, gasoline stations -0.67 versus -0.67 pp, HPC $+1.08$ versus $+1.07$ pp), confirming that non-essential closures do not alter the essential-category findings through cross-category equilibrium effects. The main text focuses on the four essential categories not because income-graded welfare loss is unique to them but because they serve daily needs for which substitution to other retail categories may be limited. Losing access to a nearby grocery store imposes a qualitatively different welfare burden than losing access to a nearby sporting goods store, even if both losses exhibit similar income gradients. The essential-category focus also reflects the policy-relevant contraction scenario, as US essential retail is undergoing sustained contraction, with pharmacy closures exceeding openings between 2018 and 2021[14] and dollar-store expansion displacing independent grocery retailers and reducing fresh-produce spending among low-income households[15,16].

Table S15. Welfare loss by NAICS retail category under simultaneous 20% contraction of all 23 categories. Categories are sorted by Q1-Q5 gap in mean R_k (most negative at top, indicating largest disadvantage for the lowest-income quintile). The four essential categories used throughout the main analysis are marked with ⋆. "Closed" reports the number of POIs removed (bottom 20% by traffic within each category). Mean R_k is averaged across the five within-metro income quintiles. Regression coefficients are from pooled OLS of welfare loss $(1 - R_k)$ on within-metro standardized log income with controls for log distance to essential retail, brand count within 15 km, percent nonwhite, percent without vehicle, and metro fixed effects ($N = 10{,}903$, HC1 robust SEs). Categories with fewer than 50 POIs (Lawn and Garden, Book Stores, Other Motor Vehicle Dealers) should be interpreted with caution due to small sample size. Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Category	Brands	Closed	Mean R_k	Gap (pp)	β_{income}	p	Sig.
Sporting Goods, Hobby, and Musical Inst.	20	133	0.862	−5.36	−0.0054	< 0.001	***
⋆ Grocery Stores	62	735	0.834	−4.50	−0.0119	< 0.001	***
Used Merchandise Stores	8	25	0.873	−4.33	−0.0015	0.183	n.s.
Furniture Stores	7	12	0.867	−2.69	−0.0032	< 0.001	***
Other Miscellaneous Store Retailers	9	62	0.917	−1.44	−0.0062	< 0.001	***
⋆ General Merchandise Stores	13	273	0.968	−1.20	−0.0026	< 0.001	***
⋆ Gasoline Stations	20	324	0.857	−0.67	−0.0024	< 0.001	***
Other Motor Vehicle Dealers	2	5	0.920	−0.34	−0.0008	0.116	n.s.
Clothing Stores	78	330	0.880	−0.27	+0.0016	< 0.001	***
Jewelry, Luggage, and Leather Goods	7	17	0.856	+0.05	−0.0047	< 0.001	***
Beer, Wine, and Liquor Stores	2	26	0.827	+0.19	+0.0054	< 0.001	***
Department Stores	11	109	0.866	+0.25	+0.0016	0.038	*
Electronics and Appliance Stores	3	27	0.799	+0.26	−0.0017	0.153	n.s.
Shoe Stores	22	95	0.821	+0.34	+0.0014	0.067	n.s.
Home Furnishings Stores	12	32	0.920	+0.44	+0.0137	< 0.001	***
Automobile Dealers	16	29	0.878	+0.52	+0.0028	< 0.001	***
Building Material and Supplies Dealers	9	142	0.990	+0.84	+0.0003	0.558	n.s.
Specialty Food Stores	10	31	0.827	+1.00	−0.0014	0.105	n.s.
⋆ Health and Personal Care Stores	21	390	0.907	+1.08	+0.0043	< 0.001	***
Auto Parts, Acc., and Tire Stores	14	174	0.900	+1.09	+0.0040	< 0.001	***
Book Stores and News Dealers	3	9	0.951	+1.12	−0.0001	0.837	n.s.
Office Supplies, Stationery, and Gift	10	55	0.849	+2.80	+0.0126	< 0.001	***
Lawn and Garden Equip. and Supplies	1	2	0.850	+3.63	+0.0081	< 0.001	***

9 Per-category and per-metro regression structure

Running the pooled regression separately for each essential category (Table S16) reveals two patterns that the aggregate conceals. The negative distance coefficient appears only for grocery and the aggregate $R_{\min}$, driven by suburban neighborhoods accessing large competitive stores despite greater distance. All other categories follow standard spatial logic, with farther neighborhoods losing more. The association with racial composition changes sign across categories. Percent nonwhite is significant and positive for grocery and general merchandise, significant and negative for HPC, and not significant in the aggregate ($p = 0.81$), indicating that the association differs by category but is not significant for the aggregate weakest-link loss measure. The share of households without a vehicle is positively associated with HPC welfare loss ($\beta = +0.086$, $p < 0.001$).

The income gradient replicates across all five metros individually (Table S17), with per-standard-deviation welfare loss ranging from -0.61 pp in San Diego to -1.94 pp in San Jose. Income is the only predictor significant in all five metros. Brand count reverses sign in Sacramento ($\beta = -0.017$, $p < 0.001$), and the association with racial composition also changes sign across metros (significant and negative in San Diego, significant and positive in San Jose).

Table S16. Per-category regression coefficients at 20% contraction. Each column is a separate OLS regression of welfare loss on neighborhood characteristics with metro fixed effects (not shown). DV $= 1 - R_k$ for each essential category and $1 - R_{\min}$ for the aggregate. $N = 10{,}903$. HC1 robust SEs in parentheses. Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

	Grocery	**HPC**	**Gasoline**	**Gen. Merch.**	$R_{\min}$
log income (z)	−0.012***	+0.004***	−0.002***	−0.003***	−0.010***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
log distance (z)	−0.006***	+0.012***	+0.008***	+0.001**	−0.007***
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
brand count (z)	+0.009***	+0.009***	−0.001	+0.001	+0.003***
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
% nonwhite	+0.019***	−0.017***	−0.000	+0.010***	+0.001
	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)
% no vehicle	−0.006	+0.086***	+0.029***	+0.012**	+0.019**
	(0.007)	(0.008)	(0.007)	(0.004)	(0.007)
R^2	0.279	0.110	0.111	0.190	0.205

Table S17. Per-metro regression coefficients at 20% contraction. Each column is a separate within-metro OLS regression with DV $= 1 - R_{\min}$ and the same five predictors (no metro fixed effects). HC1 robust SEs in parentheses. Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

	LA *n=5,714*	**SF** *n=1,814*	**SD** *n=1,475*	**SJ** *n=819*	**SAC** *n=1,081*
log income (z)	−0.007***	−0.014***	−0.006***	−0.019***	−0.014***
	(0.001)	(0.002)	(0.002)	(0.003)	(0.002)
log distance (z)	−0.007***	−0.015***	+0.002	−0.018***	−0.003
	(0.001)	(0.002)	(0.001)	(0.002)	(0.003)
brand count (z)	+0.005***	+0.004**	+0.014***	+0.002	−0.017***
	(0.001)	(0.002)	(0.001)	(0.003)	(0.003)
% nonwhite	+0.003	+0.003	−0.029***	+0.064***	−0.005
	(0.003)	(0.007)	(0.005)	(0.011)	(0.006)
% no vehicle	+0.030**	−0.009	+0.048*	−0.049	+0.048
	(0.010)	(0.013)	(0.024)	(0.034)	(0.026)
R^2	0.071	0.116	0.092	0.250	0.188

10 Robustness of the spatial-versus-demographic decomposition

Table S18 reports counterfactual gap closure under two alternative income groupings. When the comparison is broadened from income quintiles (Q1 vs Q5) to median-split halves (bottom vs top half of the within-metro income distribution), CF-A (spatial swap) closes 179% of the gap at 20% contraction, while CF-B (demographic swap) closes 13%. Gap closure exceeding 100% indicates a reversal of the original income gradient, with lower-income neighborhoods attaining higher welfare retention than higher-income neighborhoods under the counterfactual. The spatial dominance holds at every contraction level tested (10-50%) and in every metro individually (CF-A closure ranges from 96% in San Jose to 271% in San Diego). When the comparison is further narrowed to the interior of the income distribution (Q2 vs Q4, excluding both tails), CF-A still closes more than 100% of the gap in all five metros (pooled 183%), while CF-B closes at most 21% (San Francisco). The consistency across these three comparison windows, Q1-vs-Q5, median-split, and Q2-vs-Q4, makes it unlikely that the spatial dominance finding is driven by outlier neighborhoods at the extremes of the income distribution. The structural differences associated with CF-A are shown in Table S19. Under CF-A, Q1 neighborhoods inherit Q5's distance profile, which is farther from the nearest surviving grocery store (1.53 vs 1.18 km), and Q1's welfare retention improves from $R_{\min} = 0.791$ to 0.842, exceeding Q5's original 0.825. The counterfactual spatial profile also features fewer nearby essential retail closures (6.0 vs 8.6 within 5 km) and fewer surviving retail locations (24.4 vs 28.1). Furthermore, the Shapley-Owen R^2 decomposition remains stable across contraction intensities, with demographic predictors accounting for 57-63% of explained variance at all contraction levels from 10% to 50%.

Table S18. Counterfactual gap closure under alternative income groupings.

Panel A: Median-split, pooled

Contraction	Original gap (pp)	CF-A (%)	CF-B (%)
10%	0.60	472	12
20%	1.99	179	13
30%	2.17	280	12
40%	2.34	220	13
50%	2.47	178	10

Panel B: Median-split, per metro at 20%

Metro	Original gap (pp)	CF-A (%)	CF-B (%)
Los Angeles	1.49	252	10
San Francisco	3.19	114	18
San Diego	0.99	271	15
San Jose	3.78	96	11
Sacramento	2.60	136	10

Panel C: Q2-vs-Q4, per metro at 20%

Metro	Original gap (pp)	CF-A (%)	CF-B (%)
Pooled	1.68	183	14
Los Angeles	1.23	250	13
San Francisco	2.63	117	21
San Diego	1.76	198	9
San Jose	2.65	128	2
Sacramento	1.65	145	18

Table S19. Spatial counterfactual mechanism at 20% contraction. Structural characteristics of Q1 neighborhoods before and after the CF-A swap (receiving Q5's within-metro median distance profile), compared with Q5's original values.

Metric	Q1 original	Q1 under CF-A	Q5 original
Mean dist. to nearest surviving grocery (km)	1.18	1.53	1.53
Essential POIs within 5 km closed	8.6	6.0	6.0
Essential POIs within 5 km surviving	28.1	24.4	24.4
Mean top-3 brand share (essential avg.)	0.504	0.506	0.497
Mean $R_{\min}$ at 20%	0.791	0.842	0.825

11 Channel decomposition of welfare loss

Holding demographic unfamiliarity at its pre-contraction value, the change in experienced welfare decomposes additively into three channels (access, congestion, substitution; see Supplementary Note 5). Table S20 reports the full grocery decomposition across income deciles. Access loss is the dominant channel at every decile, with a monotonic income gradient from -27.5 to -23.1 log-welfare points (D1 to D10). The substitution channel changes sign between the fourth and fifth income deciles, running from -0.8 log-welfare points for the poorest decile to $+1.4$ for the richest (Spearman $\rho = +0.17$ with continuous income, $p < 0.001$). Among the lowest four deciles, substitution further reduces welfare as consumers are pushed toward less-preferred alternatives after nearby stores close. Among higher-income deciles, reallocation toward better-matched alternatives partially offsets access losses. The congestion channel contributes positively across all deciles ($+5.8$ to $+6.4$ log-welfare points), partially offsetting access losses, but varies little with income ($\rho = -0.05$). Its contribution to the Q1-Q5 grocery gap is consequently small ($+0.34$ log-welfare points) compared with those of access and substitution. The access channel also correlates positively with income ($\rho = +0.20$, $p < 0.001$), so both access and substitution reinforce the income gradient in grocery welfare loss.

This compounding is most pronounced in grocery (Table S21). In HPC, substitution contributes substantially to welfare retention for both Q1 and Q5 ($+44.2$ and $+42.8$ log-welfare points, respectively), consistent with consumers reallocating visits toward surviving major retail chains. The Q1-Q5 substitution gap in HPC is $+1.4$ log-welfare points, meaning substitution slightly favors Q1. Access loss also slightly favors Q1 ($+0.3$ log-welfare points), and the combined channel contributions produce a reversed aggregate gradient ($+1.3$ log-welfare points net, with Q5 experiencing greater losses). Gasoline stations show a similar pattern of positive substitution buffering access loss, with the net gap narrowing to -0.9 log-welfare points. General merchandise also exhibits reinforcing access and substitution disadvantages for Q1, but at smaller magnitudes (-1.0 and -0.5 log-welfare points, respectively), producing a net gap of -1.2 log-welfare points. For grocery, the corresponding mean welfare losses are 19.1% for Q1 and 14.6% for Q5, a difference of 4.5 percentage points. Grocery's distinctive substitution asymmetry, in which lower-income neighborhoods experience additional losses while higher-income neighborhoods benefit from reallocation, helps explain why it is the binding essential category for the majority of neighborhoods.

Table S20. Grocery channel decomposition by income decile at 20% contraction. Channel contributions and total welfare changes are expressed in log-welfare points ($100\Delta\widetilde{W}$), averaged across neighborhoods within each income decile. Access measures the direct utility change from store closures holding consumer behavior fixed. Congestion measures the change in the congestion component at surviving stores. Substitution measures the welfare effect of consumers reallocating visits across remaining brands. Total $\Delta\widetilde{W}$ is the sum of the three channels.

Decile	Access	Congestion	Substitution	Total $100\Delta\widetilde{W}$
D1 (lowest)	-27.51	$+6.41$	-0.81	-21.91
D2	-26.88	$+6.35$	-0.81	-21.34
D3	-26.23	$+6.28$	-0.45	-20.40
D4	-25.36	$+6.02$	-0.02	-19.36
D5	-24.92	$+6.04$	$+0.10$	-18.77
D6	-24.57	$+5.92$	$+0.42$	-18.23
D7	-24.07	$+6.00$	$+0.65$	-17.42
D8	-23.50	$+5.82$	$+1.04$	-16.64
D9	-23.41	$+5.87$	$+1.24$	-16.30
D10 (highest)	-23.14	$+6.21$	$+1.42$	-15.51

Table S21. Channel decomposition by essential category and income quintile at 20% contraction. Channel contributions and total welfare changes are expressed in log-welfare points ($100\Delta\widetilde{W}$). The gap is calculated as Q1 minus Q5 for each channel and total, with negative values indicating greater welfare losses for Q1.

Category	**Quintile**	**Access**	**Cong.**	**Subst.**	**Total** $100\Delta\widetilde{W}$
Grocery	Q1	−27.20	+6.38	−0.81	−21.62
	Q5	−23.27	+6.04	+1.33	−15.90
	Gap	−3.93	+0.34	−2.14	−5.72
HPC	Q1	−60.27	+6.90	+44.21	−9.16
	Q5	−60.55	+7.31	+42.83	−10.41
	Gap	+0.28	−0.42	+1.39	+1.25
Gasoline	Q1	−29.61	+1.46	+12.14	−16.01
	Q5	−28.08	+1.89	+11.03	−15.16
	Gap	−1.53	−0.43	+1.11	−0.85
Gen. Merch.	Q1	−11.27	+0.48	+6.92	−3.87
	Q5	−10.24	+0.25	+7.37	−2.63
	Gap	−1.03	+0.24	−0.45	−1.24

12 Vulnerability trap conditions across categories and metros

The income gradient in welfare loss is monotonic across all five income quintiles at every contraction level tested (Table S22). The Q1-Q5 gap widens from 1.11 pp at 10% contraction to 4.10 pp at 50%, indicating that deeper commercial contraction progressively concentrates welfare losses on disadvantaged neighborhoods. To examine whether this pattern satisfies the conditions associated with a vulnerability trap, Table S23 reports three operational indicators under simulated brand-preserving contraction: income-associated closure selection, gap amplification with contraction intensity, and asymmetric welfare recovery. These indicators are evaluated for the pooled $R_{\min}$ outcome under three closure strategies (Panel A) and for each essential category under all three strategies (Panel B). Under traffic-based closure, the pooled $R_{\min}$ gap peaks at 4.69 pp, accompanied by a significant income gradient in closure probability ($\beta_{\text{income}} = -0.148$, $p < 0.001$), a strong positive association between the welfare gap and contraction intensity ($\rho = +0.85$, $p = 0.001$), and recovery asymmetry of +2.9 pp favoring Q5. Competition-based closure produces a larger-magnitude negative income coefficient ($\beta = -0.270$, $p < 0.001$), a peak gap of 4.37 pp, and recovery asymmetry of +2.8 pp. Random closure generates a substantially smaller peak gap (1.45 pp) and no significant income gradient in closure probability ($\beta = -0.020$, $p = 0.42$), although the gap still increases with contraction intensity ($\rho = +0.70$, $p < 0.05$). These results suggest that underlying spatial vulnerability contributes to unequal welfare losses even under random closure, while market-driven closure orderings substantially amplify the disparity. The per-category results reveal that grocery exhibits the largest peak gap and strongest positive gap monotonicity under traffic-based closure (3.68 pp, $\rho = +0.92$, $\beta = -0.12$), although HPC and general merchandise exhibit larger-magnitude negative income coefficients. Gasoline stations provide a contrasting case. Although their income-associated closure coefficient is comparable to grocery's ($\beta = -0.12$, $p = 0.035$), the welfare gap does not exhibit significant progressive widening ($\rho = +0.44$, $p = 0.180$). This is consistent with greater substitution opportunities moderating the cumulative distributional effects of gasoline-station closures, despite measurable welfare disparities and recovery asymmetry. The full gap trajectories across all ten contraction steps are shown in Fig. S1.

Table S22. $R_{\min}$ **by income quintile across contraction levels.** Category-proportional contraction at five intensity levels.

Contraction	Q1	Q2	Q3	Q4	Q5	Gap (pp)
10%	0.899	0.904	0.905	0.909	0.910	1.11
20%	0.791	0.800	0.807	0.817	0.825	3.38
30%	0.690	0.697	0.703	0.717	0.724	3.43
40%	0.584	0.591	0.596	0.612	0.621	3.65
50%	0.482	0.489	0.494	0.510	0.523	4.10

Table S23. Trap conditions under brand-preserving contraction. Three operational indicators are evaluated under simulated closure: C1 (income-associated closure selection, measured by the logistic coefficient β_{income} of closure probability on neighborhood income), C2 (gap amplification, measured by the peak Q1-Q5 welfare gap and Spearman ρ of the signed gap with contraction step), and C3 (recovery asymmetry, measured as Q5 minus Q1 welfare recovery at 50% restoration). The signed gap is defined as $100(\overline{R}_{Q5} - \overline{R}_{Q1})$, with positive values indicating greater welfare loss for Q1. **Panel A:** Pooled $R_{\min}$ under three closure orderings. **Panel B:** Per-category results under traffic-, competition-, and random-based closure. Dashes indicate C1 coefficients not reported for category-specific competition-based closure. Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Panel A: Pooled $R_{\min}$

Strategy	C1: β_{inc}	C2: Peak (pp)	C2: ρ	C3: Rec. (pp)
Traffic	−0.148***	4.69	+0.845**	+2.9
Competition	−0.270***	4.37	+0.655*	+2.8
Random	−0.020 n.s.	1.45	+0.700*	+1.2

Panel B: Per-category, all strategies

Category	Strategy	C1: β_{inc}	C2: Peak (pp)	C2: ρ	C3: Rec. (pp)
Grocery	Traffic	−0.119**	3.68	+0.918***	+2.68
	Competition	—	4.26	+0.964***	+0.83
	Random	−0.003 n.s.	0.04	−0.891***	−0.22
HPC	Traffic	−0.228***	3.15	+0.773**	+2.14
	Competition	—	4.53	+0.918***	−0.26
	Random	−0.032 n.s.	0.36	+0.773**	−0.37
Gen. Merch.	Traffic	−0.137*	3.56	+0.900***	+2.16
	Competition	—	0.64	+0.045 n.s.	+1.39
	Random	+0.072 n.s.	0.39	−0.545 n.s.	+1.00
Gasoline	Traffic	−0.123*	2.10	+0.436 n.s.	+2.62
	Competition	—	0.00	−0.127 n.s.	+0.51
	Random	+0.027 n.s.	0.00	−0.973***	−0.12

Grocery gap amplification under traffic-based contraction is observed in three of the five metropolitan areas, although the direction and strength of the trend vary across metros (Table S24). Los Angeles shows the strongest positive gap monotonicity ($\rho = +0.91$, $p < 0.001$), with a peak Q1-Q5 gap of 6.44 pp. San Francisco ($\rho = +0.86$, $p < 0.001$) and Sacramento ($\rho = +0.70$, $p = 0.016$) also exhibit significant gap widening. San Diego instead shows a significant negative association between contraction intensity and the signed welfare gap ($\rho = -0.87$, $p < 0.001$), indicating a disparity in the opposite direction rather than progressive disadvantage for Q1 neighborhoods. San Jose shows no significant trend ($\rho = -0.44$, $p = 0.18$). These heterogeneous trajectories suggest that gap amplification depends on metropolitan retail organization and the spatial distribution of closures, rather than following inevitably from baseline income inequality. Differences in store concentration, neighborhood exposure, and available substitutes may influence whether market-driven contraction produces progressively widening welfare disparities.

Table S24. Per-metro grocery gap monotonicity under traffic-based brand-preserving contraction. Spearman correlation between contraction step (1–10) and the signed Q1-Q5 welfare gap in grocery, tested within each metro. The gap is defined as $100(\bar{R}_{Q5} - \bar{R}_{Q1})$, with positive values indicating greater welfare loss for Q1 and negative values indicating greater welfare loss for Q5. Maximum gaps are the largest signed values observed across contraction steps.

Metro	Gap ρ	p	Sig.	Max gap (pp)
Los Angeles	+0.909	< 0.001	***	6.44
San Francisco	+0.855	< 0.001	***	2.10
Sacramento	+0.700	0.016	*	4.01
San Jose	−0.436	0.180	n.s.	0.55
San Diego	−0.873	< 0.001	***	0.00

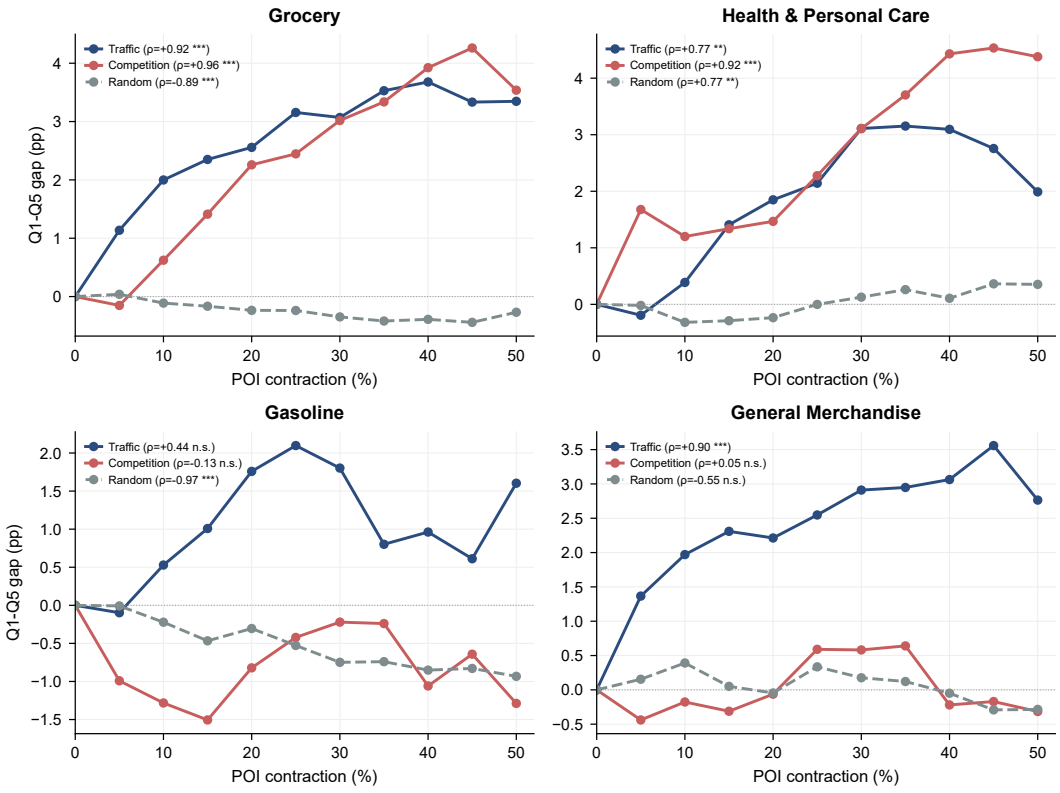

Figure S1. Q1-Q5 welfare gap trajectories under brand-preserving contraction by essential category. Each panel shows the signed Q1-Q5 gap in mean R_k (pp) across ten progressive closure steps (0–50% contraction) for three closure orderings. Positive values indicate greater welfare loss for Q1, while negative values indicate greater loss for Q5. Spearman ρ between contraction step and the signed welfare gap is annotated for each trajectory.

References

1. US Census Bureau. American Community Survey (ACS). https://census.gov/programs-surveys/acs (2024).
2. McKelvey, R. D. & Palfrey, T. R. Quantal response equilibria for normal form games. *Games economic behavior* **10**, 6–38 (1995).
3. Borkar, V. S. *Stochastic approximation: a dynamical systems viewpoint*, vol. 100 (Springer, 2008).
4. Bai, S., Kolter, J. Z. & Koltun, V. Deep equilibrium models. *Adv. neural information processing systems* **32** (2019).
5. Blondel, M. *et al.* Efficient and modular implicit differentiation. *Adv. neural information processing systems* **35**, 5230–5242 (2022).
6. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
7. Berry, S. T. Estimating discrete-choice models of product differentiation. *RAND J. Econ.* **25**, 242–262 (1994).
8. Berry, S., Levinsohn, J. & Pakes, A. Automobile prices in market equilibrium. *Econometrica* **63**, 841–890, DOI: 10.2307/2171802 (1995).
9. McFadden, D. Econometric models of probabilistic choice. In Manski, C. F. & McFadden, D. (eds.) *Structural Analysis of Discrete Data with Econometric Applications*, 198–272 (MIT Press, Cambridge, MA, 1981).
10. Anderson, S. P., de Palma, A. & Thisse, J.-F. *Discrete Choice Theory of Product Differentiation* (MIT Press, Cambridge, MA, 1992).
11. McFadden, D. Modelling the choice of residential location. In Karlqvist, A., Snickars, F. & Weibull, J. (eds.) *Spatial Interaction Theory and Planning Models*, 75–96 (North-Holland, Amsterdam, 1978).
12. Ben-Akiva, M. E. & Lerman, S. R. *Discrete choice analysis: theory and application to travel demand*, vol. 9 (MIT press, 1985).
13. Deaton, A. & Muellbauer, J. *Economics and consumer behavior* (Cambridge university press, 1980).
14. Guadamuz, J. S., Alexander, G. C., Kanter, G. P. & Qato, D. M. More US pharmacies closed than opened in 2018–21; independent pharmacies, those in Black, Latinx communities most at risk. *Heal. Aff.* **43**, 1703–1711, DOI: 10.1377/hlthaff.2024.00192 (2024).
15. Lopez, R. A., Marchesi, K. & Steinbach, S. Dollar store expansion and independent grocery retailer contraction. *Appl. Econ. Perspectives Policy* **46**, 514–533 (2024).
16. Caoui, E. H., Hollenbeck, B., Osborne, M. & Page, E. The impact of dollar store expansion on local market structure and food access. *Quant. Mark. Econ.* **24**, 7 (2026).