Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Urban retail contraction reveals a spatial vulnerability trap in essential retail access

Xinwu Qian

xq15@rice.edu

Rice University https://orcid.org/0000-0001-8001-2164

Hossein Gazmeh

Rice University https://orcid.org/0009-0007-1008-7225

Article

Keywords:

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11652820/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

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, United States
[*]xq15@rice.edu

ABSTRACT

Grocery stores, pharmacies, and other essential retailers anchor daily life, but in U.S. cities they are closing unevenly. These closures do more than remove nearby options. They redirect customer flows through urban retail networks and redistribute welfare losses across neighborhoods. To measure this reorganization, we develop and validate a differentiable equilibrium model of store choice that estimates visit and welfare shifts after closures, using GPS mobility data linking 11,252 neighborhoods to 360 retail brands across five California metros. Simulating closure of the least-visited 20% of essential retail locations, we find welfare losses are 3.4 percentage points higher for the lowest-income quintile than for the highest, with a monotonic gradient in all five metros. Counterfactuals giving low-income neighborhoods the demographics of high-income ones, while preserving retail geography, close at most 18% of this gap. The core mechanism is concentrated dependence on nearby, hard-to-replace essential stores. Market-driven contraction amplifies this exposure and favors advantaged neighborhoods during recovery, characterizing a spatial vulnerability trap. Breaking this trap requires identifying and stabilizing exposed essential retail before contraction. These results suggest that proximity to stores can mask dependence on them, and that modeling how people reallocate visits can flag at-risk neighborhoods before closures occur.

Introduction

Essential retail in U.S. cities is undergoing sustained and uneven contraction. More pharmacies closed than opened between 2018 and 2021, with closures falling disproportionately on predominantly Black and Latinx neighborhoods[1]. Grocery consolidation is following a similar trajectory, as dollar-store expansion displaces independent retailers and reduces fresh-produce spending among low-income households[2,3]. Roughly 18% of Americans live in "pharmacy deserts," where access is limited by travel time, while another 9% rely on a single accessible pharmacy[4]. Because grocery, pharmacy, and other essential retailers anchor everyday urban welfare, this contraction raises a prospective urban policy question: which neighborhoods would lose the most welfare under future closures, and how could those losses be anticipated before they occur? Existing store maps and visit data record today's retail system. They do not identify how households would substitute after closures, how remaining stores would absorb redirected demand, or where welfare would fall once the retail network reorganizes. We address this question with a mobility-calibrated substitution equilibrium that identifies where welfare losses concentrate and explains why market recovery can bypass the neighborhoods most exposed to contraction.

The standard framework for asking who is harmed by retail contraction is proximity-based accessibility. Neighborhoods far from essential retail are labeled food or pharmacy deserts, while neighborhoods with nearby stores are counted as adequately served[5–7]. Proximity captures a real dimension of disadvantage, but it answers a static question of who is close to existing stores. Retail contraction poses a counterfactual question of who can substitute effectively after stores disappear. A closure changes the choices available to neighborhoods. It removes a destination, shifts consumers toward surviving alternatives, and changes the attractiveness of those alternatives as demand reallocates. The neighborhoods most harmed after contraction can therefore differ from those that appear least served before contraction.

This counterfactual lens makes commercial dependence the central focus of our study. In dense urban retail systems, a neighborhood can appear well served because essential stores are nearby, while its residents depend on a narrow set of stores that are difficult to replace. Welfare loss from a closure depends on the quality of feasible substitutes, the congestion created at those substitutes, and the fit between surviving stores and residents' routine shopping patterns. These substitution capacities are uneven. Low-income neighborhoods often face thinner retail assortments, higher per-unit prices, longer travel burdens, limited vehicle access, and greater exposure to lower-quality outlets[5,6,8,9]. Under contraction, proximity can conceal dependence, as the stores that make a neighborhood appear served may be precisely those whose loss is hardest to absorb.

Existing research identifies important pieces of this process while leaving the counterfactual reorganization unmeasured. Access studies document where essential retailers are located and how these patterns correlate with income and race, but they describe the pre-closure landscape rather than the welfare losses that emerge after disruption[4–7]. Mobility-based studies show

that residential buffers understate realized disadvantage because activity spaces extend beyond neighborhoods and because travel time, operating hours, and vehicle access shape where people shop[10–13]. These studies do not recover the equilibrium response to future closures. Network models capture business interdependencies without micro-founded welfare measures[14,15], while structural demand models recover welfare under changing market conditions but do not directly resolve neighborhood-level visit reallocation and congestion following retail closures[3,16,17]. Assessments of retail vulnerability therefore often rely on static proxies, such as nearby store counts or distance to the next retailer[4]. The missing object is a neighborhood-resolution substitution equilibrium with congestion feedback and welfare predictions, calibrated against observed consumer movement while remaining differentiable enough to evaluate counterfactual closures.

We build this object as a differentiable model of consumer retail choice trained on GPS mobility data covering approximately 15.8 million recorded visits per month across 360 retail brands, 11,252 census block groups (CBGs), and five California metropolitan areas. The model converts an observed store-visit network into the post-closure welfare distribution implied by any given contraction. We assess its counterfactual predictions by comparing simulated visit redistribution after Rite Aid's removal with observed changes in retail traffic following its 2023 bankruptcy. Applying the model to progressive essential retail contraction and subsequent market-driven recovery reveals a spatial "vulnerability trap" in which low-income neighborhoods systematically lose more welfare, the gap widens through roughly 30% contraction, and partial recovery favors higher-income neighborhoods. Structural counterfactual experiments identify spatial exposure, rather than demographic composition, as the dominant mechanism within the model. Contraction compounds the commercial dependence that made these neighborhoods vulnerable, and partial recovery does not immediately undo the resulting disparities. Breaking this cycle requires identifying and stabilizing exposed essential retail before contraction, rather than relying only on proximity-based expansion or income-targeted policy.

Results

The model captures neighborhood choice dynamics and predicts observed redistribution

To recover how neighborhoods substitute after retail closures, we develop Differentiable Brand Ecosystem Dynamics (DBED), a mobility-calibrated equilibrium framework for neighborhood retail choice with endogenous brand-attractiveness feedback. The model takes as input monthly visit records linking approximately 15.8 million observed visits from 11,252 CBGs to 360 retail brands across 23 retail categories in five California metropolitan areas (Fig. 1a). For each CBG-brand pair, utility combines six interpretable features (distance, income-affordability match, brand size, dwell time, online transaction share, and demographic unfamiliarity) with a neural residual that captures nonlinear interactions (Fig. 1b). The stationary visit shares are then determined as the quantal response equilibrium solved through fixed-point iteration. As consumers reallocate across stores, brand attractiveness adjusts, and visit shares update until they are consistent with the resulting utilities. The model is trained end-to-end by differentiating through a single equilibrium step from the converged fixed point, enabling gradient-based learning without storing the fixed-point iterations. We focus the welfare analysis on four essential retail categories that anchor daily life for urban households, namely Grocery Stores, Health and Personal Care Stores (HPC, including drugstores and pharmacies), Gasoline Stations, and General Merchandise Stores (including warehouse clubs and supercenters). See Methods for the model specification and training procedure.

On held-out CBGs, DBED predicts observed retail demand with a Spearman correlation (ρ) of 0.73 on visit counts ($R^2 = 0.62$) and 0.60 on within-category visit shares (Fig. 1c). This performance improves on a gravity baseline with grid-searched distance-decay parameter of $\alpha = 1.25$ ($\rho = 0.44$ on visits and 0.41 on shares; Supplementary Note 7) and on a uniform-allocation baseline ($\rho = 0.32$, $R^2 = 0.17$ on visits; $\rho = 0.14$ on shares). The learned coefficients also have theoretically consistent signs (Fig. 1d). Distance enters negatively and most strongly ($\beta_{\text{dist}} = -0.97$), while affordability match enters positively ($\beta_{\text{afford}} = +0.47$). Demographic unfamiliarity enters negatively ($\beta_{\text{unfam}} = -0.47$), indicating lower visit propensity when a brand's visitor base is demographically less similar to the origin neighborhood. Brand size ($\beta_{\text{size}} = +0.32$) and dwell time ($\beta_{\text{dwell}} = +0.23$) increase utility, while online transaction share enters weakly negatively ($\beta_{\text{online}} = -0.07$), consistent with mild substitution between digital and physical retail activity. The learned mixing weight places a weight of 0.63 on the interpretable component and 0.37 on the neural residual, making the model structurally transparent while retaining flexibility to capture unobserved preference heterogeneity. Predictive accuracy on held-out CBGs, however, is only a first requirement for counterfactual use. A model that interpolates well across neighborhoods may still fail to predict how visits redistribute when a brand exits a market.

The 2023 Rite Aid bankruptcy and subsequent California store closures[18] provide an opportunity to assess counterfactual redistribution against observed retail traffic changes. Trained on 2022–2023 mobility data and given no information about the observed 2024 traffic changes, DBED predicts which HPC brands absorb displaced traffic when Rite Aid is removed from the choice set. Predicted brand-level traffic gains correlate with observed 2024 traffic changes at $\rho = 0.83$ (Spearman) and $r = 0.88$ (Pearson on log scale) across the 16 HPC brands with positive observed gains (Fig. 1e). Across all 20 HPC brands, including four brands whose traffic declined during the same period, the full-sample Spearman correlation is 0.55. The model

underestimates absolute traffic magnitudes, as observed 2024 gains may also reflect broader market growth and data-panel expansion, whereas the counterfactual removes only Rite Aid. The top four predicted beneficiaries (CVS, Walgreens, ULTA Beauty, and Bath & Body Works) match the top four observed gainers and together account for 56% of displaced Rite Aid traffic in the counterfactual. These results show that DBED recovers the direction and concentration of post-closure substitution, providing empirical support for its use in the counterfactual contraction simulations below.

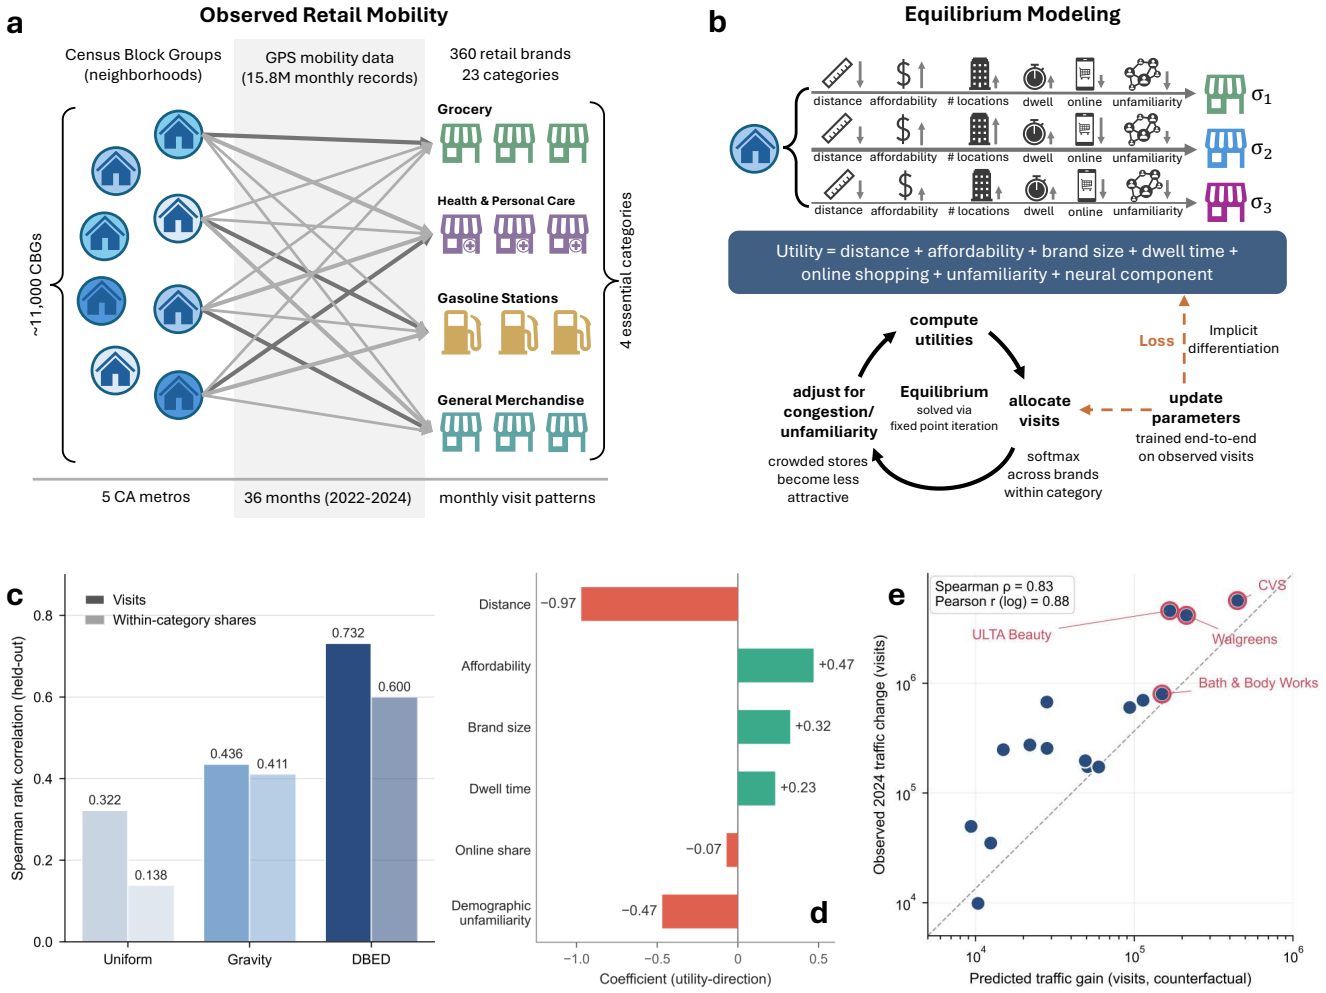

Figure 1. Differentiable equilibrium model of consumer retail choice and validation. a, Mobility-derived data structure linking 11,252 census block groups (CBGs) to 360 retail brands across 23 categories. **b,** Differentiable Brand Ecosystem Dynamics (DBED) framework for demand-responsive retail choice equilibrium. **c,** Held-out predictive performance on visit counts and within-category visit shares, compared with uniform-allocation and gravity baselines. **d,** Learned interpretable utility coefficients; positive values increase predicted utility and negative values decrease predicted utility. **e,** Counterfactual validation using the Rite Aid closures. Predicted brand-level traffic gains after removing Rite Aid from the choice set are compared with observed 2024 traffic changes across Health and Personal Care brands. Reported correlations are calculated across the 16 brands with positive observed traffic gains.

Two additional tests confirm the model captures the local competitive structure that subsequent sections rely on to generate differential welfare outcomes. CBGs with no HPC alternative within 3 km of the nearest Rite Aid show 34% higher equilibrium dependence on Rite Aid than CBGs with at least one alternative (0.069 vs 0.051 share, $n = 1{,}595$ CBGs without nearby alternatives vs 9,657 with alternatives, $p < 0.001$). This result shows that the model assigns greater dependence on Rite Aid to neighborhoods with limited nearby competition. A complementary test distinguishes geographic concentration of dependence from diffuse market share. Chevron, Shell, ARCO, and Target each command mean within-category shares of 15% to 27% across nearly all CBGs, far larger than Rite Aid's 5.6% mean share concentrated in 725 CBGs. Removing these large-share brands redirects their traffic to a few dominant alternatives, with a substitution Herfindahl-Hirschman Index (HHI) of 0.16–0.21, whereas Rite Aid's traffic disperses across many HPC competitors (HHI 0.118). These comparisons distinguish local dependence on a brand from its aggregate market share, a distinction that matters because the neighborhoods most harmed by contraction are not necessarily those where the affected brand had the largest footprint. The Rite Aid validation and these

two tests support using the model to address the central question of how commercial contraction redistributes welfare across neighborhoods.

Commercial contraction produces structurally unequal welfare losses

Essential retail welfare losses decline monotonically with neighborhood income under simulated commercial contraction, even though lower-income neighborhoods are closer to essential retail in the metros we study. We simulate contraction by progressively removing the lowest-traffic points of interest within each essential category at five levels (10% through 50%) and re-solving the full consumer choice equilibrium at each level (see Counterfactual simulations). For each neighborhood, we measure welfare retention in category k as $R_k = \exp(\Delta W_k)$, where ΔW_k is the change in share-weighted, congestion-adjusted utility relative to that neighborhood's own pre-contraction baseline (see Welfare measurement). Values below one indicate welfare loss through some combination of increased distance to surviving stores, congestion at those stores, and forced substitution to less familiar alternatives. Because each neighborhood is compared to its own baseline, R_k measures changes in model-implied experienced utility on an exponentiated scale, independent of absolute baseline utility. Overall essential welfare retention is captured by $R_{\min} = \min_k R_k$ across the four essential categories, reflecting the weakest link in a neighborhood's essential retail ecosystem. At 20% contraction, mean $R_{\min}$ declines from 0.825 for the highest within-metro income quintile (Q5) to 0.791 for the lowest (Q1), a gradient that is strictly monotonic across all five quintiles and replicates at every contraction level from 10% to 50% (Supplementary Note 12). The Q1-Q5 gap widens from 1.1 percentage points at 10% contraction to 4.1 pp at 50%. At 20% contraction, 46.6% of Q1 neighborhoods exceed 20% welfare loss, compared with 34.0% of Q3 and 21.6% of Q5 (Fig. 2b), confirming that the disparity extends beyond mean differences to the concentration of large welfare losses.

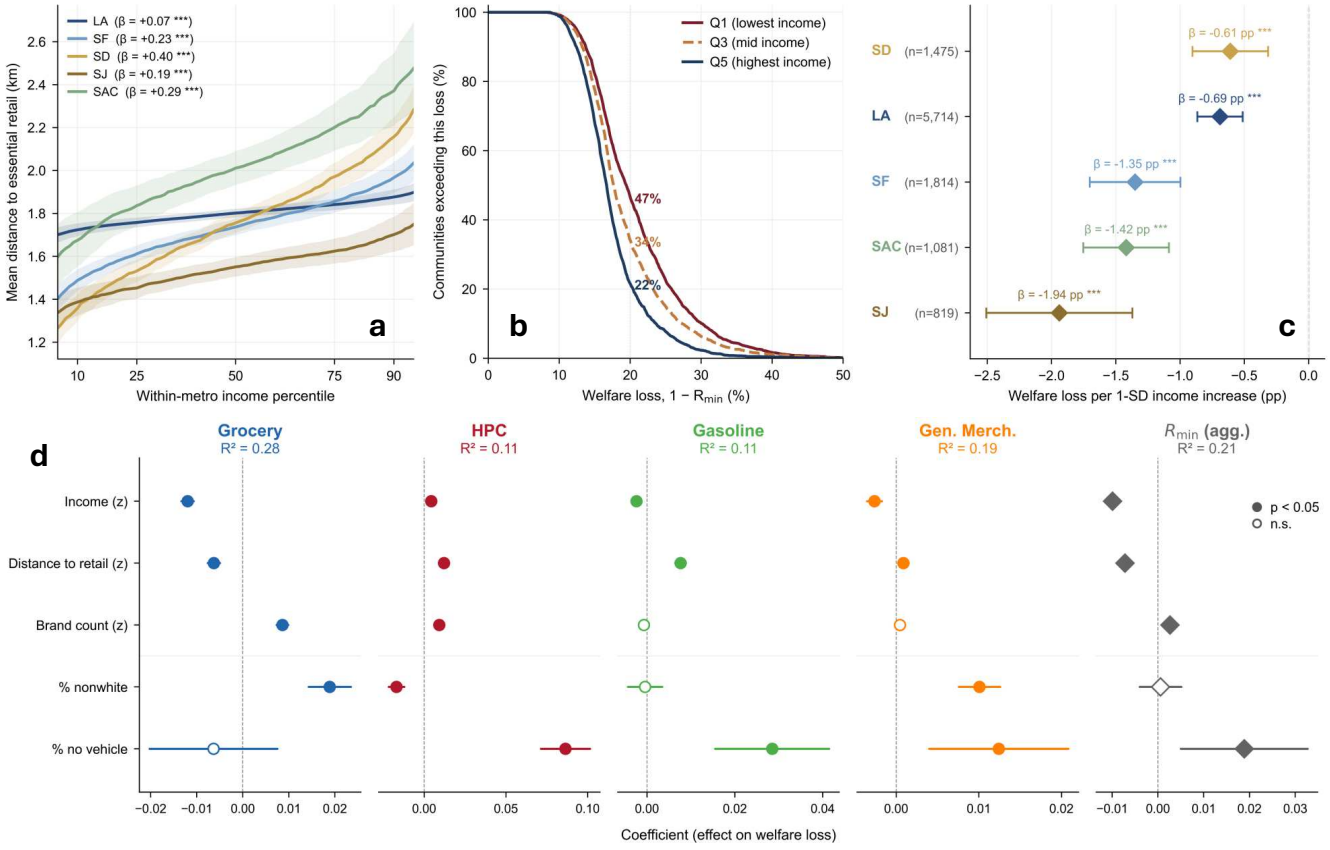

Figure 2. Commercial contraction produces structurally unequal welfare losses graded by neighborhood income. a, Mean distance to essential retail (average of nearest-brand distance) increases with income in all five metros. β denotes the log-log OLS slope. **b,** Exceedance curves of welfare loss ($1 - R_{\min}$) at 20% category-proportional contraction for income quintiles Q1 (lowest), Q3 (middle), and Q5 (highest). Percentages indicate the share of neighborhoods exceeding 20% welfare loss. **c,** Per-metro OLS regression coefficients showing welfare loss per one-standard-deviation increase in within-metro log income (points and 95% confidence intervals from within-metro regressions with spatial controls). N indicates the number of census block groups per metro. **d,** OLS regression coefficients from identical specifications (welfare loss on neighborhood characteristics with metro fixed effects) run separately for each essential category and for $R_{\min}$ (aggregate). $N = 10{,}903$ for all regressions.

We next show that proximity alone does not explain this disparity. In all five metropolitan areas, low-income neighborhoods

are closer to essential retail than high-income neighborhoods (Fig. 2a), with Q1 averaging 1.4-1.8 km to the nearest essential retailer compared with 1.8-2.4 km for Q5 depending on the metro. Nevertheless, this proximity does not protect them when that retail contracts. A pooled regression of welfare loss on neighborhood characteristics at 20% contraction ($N = 10{,}903$, $R^2 = 0.205$) reveals a significant negative association with income ($\beta = -0.010$, $p < 0.001$), with a one-standard-deviation decrease in within-metro log income corresponding to 1.0 pp of additional welfare loss (Fig. 2d). The number of essential brands within 15 km is positively associated with welfare loss ($\beta = +0.003$, $p < 0.001$), suggesting that a greater number of nearby brands does not necessarily buffer neighborhoods against contraction. Grocery is the binding essential category, determining $R_{\min}$ for 55% of neighborhoods at 20% contraction. Among the four essential categories, grocery shows the largest income coefficient ($\beta = -0.012$, $R^2 = 0.28$), with Q1 losing 19.1% of grocery welfare versus 14.6% for Q5 (a 4.5 pp difference). Gasoline stations are near-equitable (0.7 pp gap), consistent with their dense spatial coverage and high substitutability.

Per-category regressions reveal patterns that the aggregate conceals (Fig. 2d). The negative distance coefficient, in which neighborhoods farther from retail lose less, appears only for grocery ($\beta = -0.006$) and the aggregate $R_{\min}$ ($\beta = -0.007$). All other categories follow standard spatial logic, with farther neighborhoods losing more. The counterintuitive aggregate pattern is driven by grocery, the binding category, consistent with suburban neighborhoods accessing large competitive stores that provide high utility despite greater distance. HPC is the only category where the income gradient reverses ($\beta = +0.004$, $p < 0.001$), with Q5 losing 1.1 pp more than Q1. This reversal may reflect the spatial distribution of low-traffic pharmacy locations in wealthier suburban areas. Neighborhoods with higher shares of households without a vehicle also experience greater losses in this category ($\beta_{\text{no_vehicle}} = +0.086$). The association with racial composition changes sign across categories. Percent nonwhite is significant and positive for grocery ($\beta = +0.019$, $p < 0.001$) and general merchandise ($\beta = +0.010$, $p < 0.001$), significant and negative for HPC ($\beta = -0.017$, $p < 0.001$), and not significant in the aggregate ($p = 0.81$). The association between racial composition and welfare loss is therefore category-specific, with no significant association in the aggregate. The income gradient replicates across all five metros individually, with per-standard-deviation welfare loss ranging from -0.61 pp in San Diego to -1.94 pp in San Jose (Fig. 2c; Supplementary Note 9). Extending to all 23 retail categories confirms that income-graded welfare loss is not unique to essential retail (7 of 23 show significant negative income coefficients; Supplementary Note 8). Income and spatial variables jointly predict which neighborhoods experience the largest losses, indicating that these impacts are shaped by both socioeconomic position and geographic exposure. Distinguishing these two dimensions is therefore necessary to understand the mechanisms through which welfare losses arise.

The income gradient in welfare loss operates through spatial exposure

A standard variance decomposition initially suggests demographic dominance. A Shapley-Owen R^2 decomposition, which attributes explained variance to variable groups by averaging marginal contributions across all possible sub-models, assigns 61% of explained variance in welfare loss at 20% contraction to demographic variables (income, race, vehicle access), versus 39% to spatial variables (distance, brand count). This split is stable across contraction levels (57-63% demographic; Supplementary Note 10). A straightforward reading would attribute commercial vulnerability, measured here as welfare loss under contraction, primarily to the socioeconomic characteristics of residents. However, cross-sectional predictive power need not reflect causal importance when the predictor variables are entangled. Income shapes where households live, and where they live shapes their spatial exposure to contraction. We use model-based counterfactual experiments to distinguish cross-sectional prediction from the relative influence of spatial and demographic inputs, following the counterfactual design that separated supply-side and demand-side sources of nutritional inequality[5].

For the channel analysis, we hold demographic unfamiliarity at its pre-contraction value and decompose the resulting change in experienced welfare into three additive channels of access, congestion, and substitution (see Welfare measurement). Access loss measures the change in available welfare holding consumer behavior fixed, capturing both reduced brand accessibility and increased distance to surviving stores. Congestion measures the crowding effect at surviving stores as displaced consumers arrive. Substitution measures the welfare effect of consumers reallocating visits across remaining brands. Across the four essential categories, these channels operate in fundamentally different ways (Fig. 3a). Grocery shows the largest combined disadvantage from access and substitution for low-income neighborhoods, producing a net Q1-Q5 gap of -5.7 log-welfare points (negative values indicate larger losses for Q1), far larger than any other category. Access loss accounts for -3.9 log-welfare points (68% of the net gap, with congestion offsetting $+0.3$ log-welfare points), but substitution contributes an additional -2.1 log-welfare points as low-income neighborhoods shift visits toward lower-utility alternatives when nearby stores close. In HPC, the pattern reverses (net gap $+1.3$ log-welfare points, Q5 loses more), driven mainly by positive substitution ($+1.4$ log-welfare points), consistent with Q1 neighborhoods redirecting visits toward surviving HPC brands. Access loss also slightly favors Q1 in this category ($+0.3$ log-welfare points). Gasoline stations show partial offsetting, with access loss (-1.5 log-welfare points) largely buffered by positive substitution ($+1.1$ log-welfare points). General merchandise follows the standard pattern at smaller magnitude (-1.2 log-welfare points net). Congestion contributes at most 0.4 log-welfare points to the Q1-Q5 gap across categories, indicating a limited role in the income gradient.

The grocery substitution channel, which is the most regressive, exhibits a monotonic gradient across the full income distribution (Fig. 3b). The substitution effect runs from -0.8 log-welfare points for the poorest decile to $+1.4$ for the richest, changing sign between the fourth and fifth income deciles. Among the lowest four deciles, substitution further reduces welfare as consumers shift toward less-preferred alternatives after nearby stores close. Among higher-income deciles, reallocation toward better-matched alternatives partially offsets access losses. Access loss follows a parallel gradient, while congestion is effectively flat. The two dominant channels therefore tilt in the same direction for grocery, compounding rather than offsetting each other. This compounding is strongest in grocery and helps explain why it is the binding essential category for the majority of neighborhoods. Full channel decompositions by income decile and by essential category are reported in Supplementary Note 11.

To compare the influence of spatial exposure and demographic composition, we conduct model-based counterfactual experiments that modify inputs to the estimated model and re-solve the full consumer choice equilibrium from scratch, so that downstream effects through congestion and substitution are captured endogenously (see Counterfactual simulations). In the spatial counterfactual (CF-A), each lowest-income-quintile neighborhood receives the within-metro median distance vector of highest-income-quintile neighborhoods to all brands, including the corresponding accessibility scaling at each contraction level, while retaining its own income, racial composition, and demographic comfort zone. In the demographic counterfactual (CF-B), each lowest-income-quintile neighborhood receives the within-metro median income, racial composition, and demographic comfort-zone parameters of the highest-income quintile while keeping its own spatial position. Across all five metropolitan areas at 20% contraction, equalizing spatial access closes at least 98% of the Q1-Q5 welfare gap, fully eliminating and reversing it in four of five metros so that Q1 neighborhoods surpass Q5 in welfare (Fig. 3c). Equalizing demographics closes at most 18% (San Francisco) and as little as 5% (Los Angeles). Q1 neighborhoods face 2.6 more essential retail closures within 5 km than Q5 (8.6 vs 6.0 closed points of interest at 20% contraction). Under CF-A, this modeled exposure differential is eliminated by construction. Q1's welfare retention ($R_{\min} = 0.842$) exceeds the original Q5 value ($R_{\min} = 0.825$). The finding is not an artifact of comparing distributional tails. When the comparison is broadened to the bottom and top halves of the within-metro income distribution, CF-A still closes 179% of the gap while CF-B closes only 13% (Supplementary Note 10). These results identify spatial exposure as the dominant influence on welfare loss within the specified model counterfactuals. These results, however, rest on a single closure rule applied at fixed contraction levels. The market's closure logic may interact with the spatial vulnerability just identified.

Market-driven contraction creates a commercial vulnerability trap

The spatial vulnerability identified above may develop into a commercial "trap" when market-driven closures amplify existing disadvantage and subsequent recovery favors less affected neighborhoods. The poverty trap literature characterizes such dynamics as mechanisms in which disadvantage perpetuates itself[19,20]. We operationalize this concept through three testable conditions, namely market selection feedback, in which closure patterns disproportionately target disadvantaged neighborhoods; gap widening acceleration, in which inequality increases as contraction deepens; and recovery asymmetry, in which rebounds favor the advantaged first. We test whether commercial contraction generates an analogous trap in retail access. To isolate the spatial mechanism from brand composition effects, we employ brand-preserving contraction in which every brand loses 50% of its points of interest over ten progressive steps (5% per step). No brand disappears entirely, so the findings cannot be driven by selective elimination of discount or low-income-serving chains (see Counterfactual simulations). Three closure orderings provide the necessary variation. Traffic-based closure removes the lowest-traffic locations first, approximating market-driven contraction. Competition-based closure removes the most spatially redundant locations first, approximating strategic consolidation. Random closure serves as a null benchmark.

Market-driven closure amplifies the structural vulnerability approximately threefold. Under traffic-based closure, the welfare gap between the lowest and highest income quintiles widens monotonically from zero to a peak of 4.69 pp at 30% contraction (Spearman $\rho = +0.85$, $p = 0.001$; Fig. 3e), consistent with the gap-widening condition. Competition-based closure produces a comparable trajectory (peak 4.37 pp, $\rho = +0.66$, $p = 0.03$). Random closure generates a gap roughly three times smaller (peak 1.45 pp). To put the magnitude in context, gaps of 3-4 pp in $R_{\min}$ are comparable to the gap under category-proportional contraction at 20% (earlier section). Under that earlier 20% category-proportional contraction, approximately 47% of Q1 neighborhoods exceed 20% welfare loss, compared with 22% of Q5 neighborhoods, a 25 pp difference in the concentration of large welfare losses. Beyond 30% contraction, the gap plateaus and slightly narrows as losses become severe enough to affect even advantaged neighborhoods (Fig. 3e). Even random contraction produces a modest welfare gap, reflecting the baseline structural vulnerability identified in the previous section. Despite their proximity to essential retail, low-income neighborhoods remain more exposed under all three closure orderings. Market-driven contraction amplifies this structural vulnerability because the market's location-selection logic, closing lower-traffic stores, disproportionately removes the locations that disadvantaged neighborhoods depend on most.

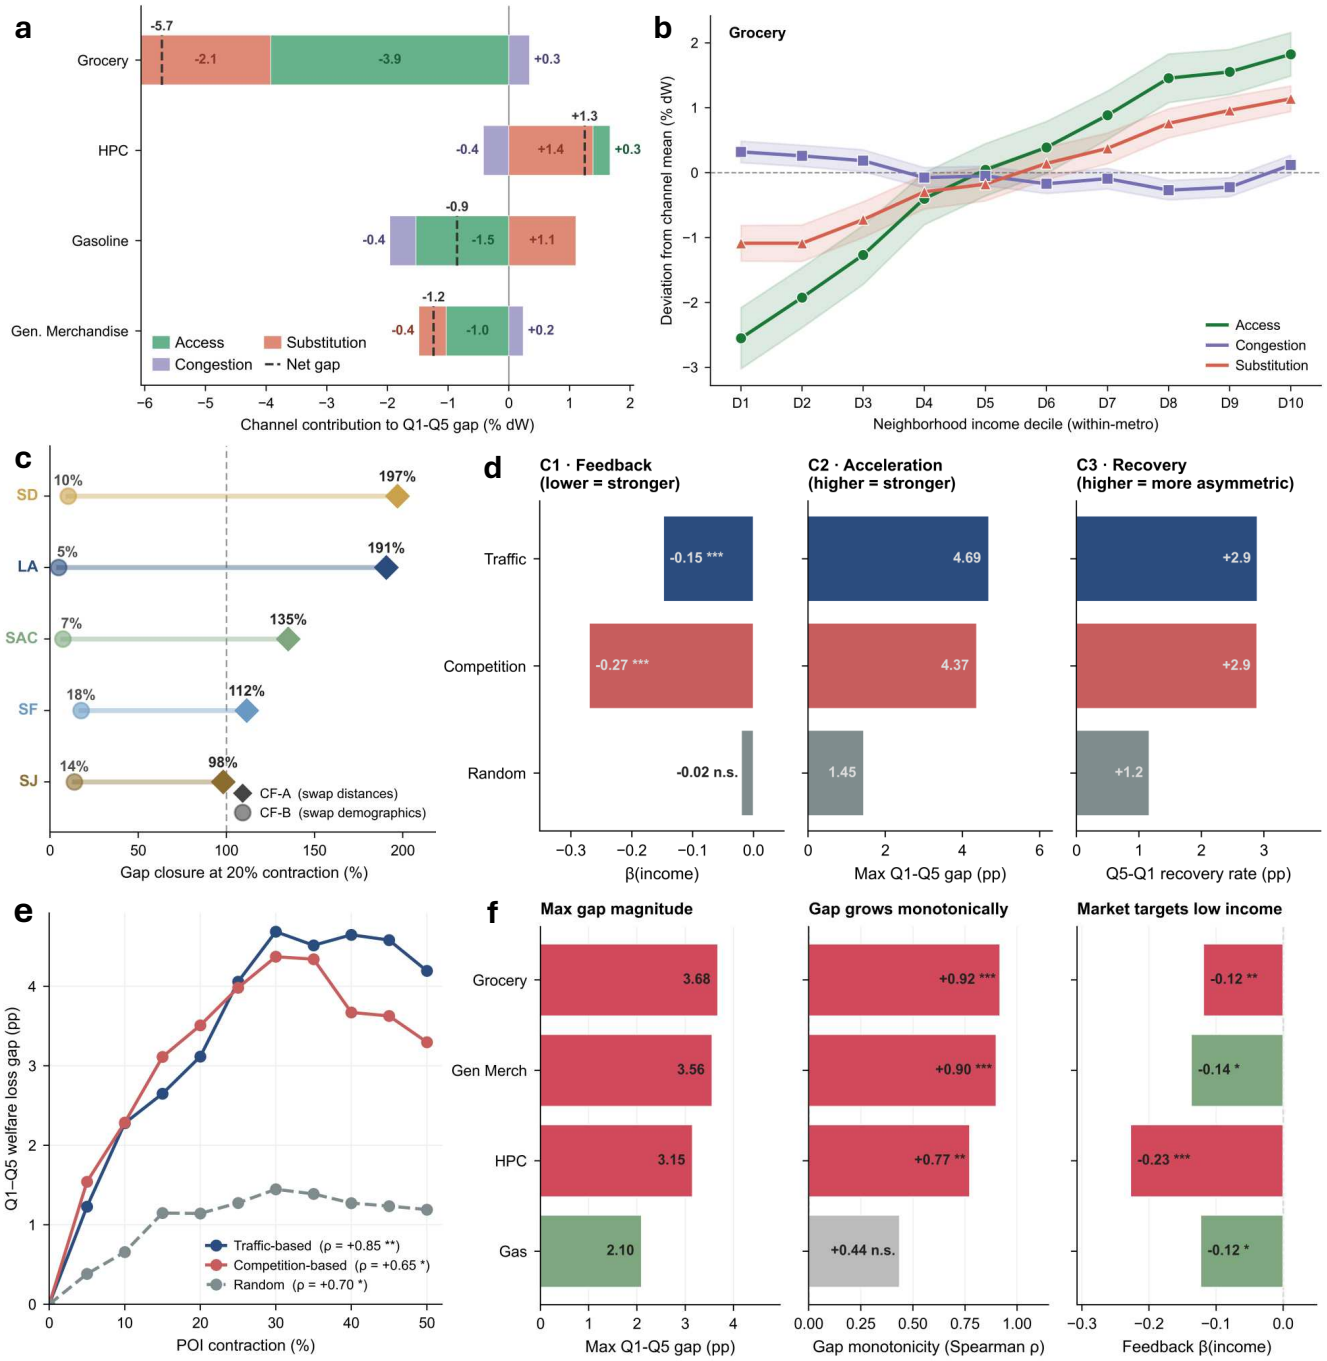

Figure 3. Spatial exposure drives welfare inequality through a self-reinforcing vulnerability trap. a, Three-channel decomposition of the Q1-Q5 welfare gap at 20% contraction by essential category. **b,** Grocery channel gradients across within-metro income deciles at 20% contraction. Each channel is mean-centered to show the income gradient on a common scale. **c,** Model-based counterfactual experiments at 20% contraction. **d,** Three trap conditions tested by closure strategy. C1: income-associated closure selection (β_{income} from logistic regression of closure on neighborhood income). C2: maximum Q1-Q5 gap across contraction steps. C3: differential recovery rate (Q5 minus Q1 welfare recovery at 50% restoration). **e,** Q1-Q5 welfare gap under progressive brand-preserving contraction for three closure orderings. **f,** Trap conditions by essential category under traffic-based closure.

The market-selection and recovery conditions confirm that this amplification is systematic (Fig. 3d). A logistic regression of closure probability on surrounding neighborhood income yields $\beta_{\text{income}} = -0.15$ ($p < 0.001$) for traffic-based and -0.27 ($p < 0.001$) for competition-based closure, indicating that each standard-deviation increase in neighborhood income lowers the odds of a nearby store closing by 14% and 24%, respectively. This income gradient in closure probability operates continuously across the income distribution, not only at the tails. Under random closure, the income coefficient is not significant ($\beta = -0.02$, $p = 0.42$), indicating that the unequal closure exposure arises from market-oriented location selection rather than random

removal. Recovery is also asymmetric. When half of the closed points of interest have reopened under traffic-based closure, Q5 neighborhoods have recovered 2.9 pp more welfare than Q1 neighborhoods, compared with 1.2 pp under random. Under the simulated recovery rule, the highest-traffic closed locations reopen first, favoring neighborhoods served by those locations. Aggregate supply can therefore rebound while the spatial gap persists or widens.

The trap operates with heterogeneous strength across essential retail categories (Fig. 3f). Grocery exhibits the strongest trap, with the largest welfare gap (3.68 pp), the strongest gap monotonicity ($\rho = +0.92$, $p < 0.001$), and significant market selection ($\beta = -0.12$, $p = 0.002$). Grocery is also the binding welfare constraint for 55% of neighborhoods and the category where access loss and substitution jointly drive the income gradient, making it the essential category most susceptible to trap dynamics across three lines of evidence. General merchandise and HPC follow similar patterns with slightly smaller gaps and weaker monotonicity but stronger income-skewed closure selection. Gasoline stations provide a contrasting case, in which closures target low-income areas but the welfare gap does not widen significantly. The closure-selection coefficient for gasoline stations ($\beta = -0.12$, $p = 0.035$) is nearly identical to grocery's, meaning the market targets low-income gasoline station locations as strongly as low-income grocery locations. However, the welfare gap does not widen significantly ($\rho = +0.44$, $p = 0.18$), consistent with gasoline stations being sufficiently ubiquitous and substitutable that losing one nearby station has limited consequences for the widening of welfare disparities. The trap thus requires not only that the market targets disadvantaged areas, which it does across all categories, but also that the targeted retail serves concentrated, hard-to-substitute local demand. Grocery exhibits all three patterns, while gasoline stations show income-skewed closure and unequal recovery without significant gap widening. The grocery gap widens significantly in three of the five metros, while San Diego exhibits a significant trend in the opposite direction and San Jose shows no significant trend (Supplementary Note 12). Full trap conditions across all four categories and three closure strategies, per-metro grocery gap monotonicity, and gap divergence trajectories are reported in Supplementary Note 12. These results indicate that the vulnerability trap is not a universal consequence of commercial contraction but emerges from the interaction between market closure logic and the spatial concentration of essential access, with the strongest dynamics in the category most central to everyday welfare.

Discussion

In the five metropolitan areas we study, low-income neighborhoods are closer to essential retail than high-income neighborhoods, so the desert framework's proximity-based logic can make them appear adequately served. Our simulations indicate that these same neighborhoods are among those most exposed to contraction. Dense nearby retail means that more locations are removed near them when lower-traffic stores close, and limited vehicle access may further constrain substitution to more distant alternatives. This finding inverts the logic of proximity-based access measurement for urban settings. Some neighborhoods that appear well served by conventional metrics are among those whose concentrated commercial dependence makes them most vulnerable when the ecosystem contracts. This vulnerability is not static. When contraction removes the lowest-traffic locations first, closures fall disproportionately in disadvantaged areas, and recovery that reinstates the highest-traffic locations first favors advantaged areas. Aggregate supply can therefore rebound while the spatial gap between neighborhoods persists or widens, so policies that monitor only total store counts would miss the distributional consequences of how contraction and recovery unfold. If the structural contraction trends that motivated this analysis continue, these losses may accumulate. The compounding nature of the trap suggests that neighborhoods that have already absorbed recent pharmacy and grocery closures would enter each successive wave from a weakened position, with fewer substitutes, longer distances, and narrower choice sets.

The policy implication follows directly from the mechanism. Expanding retail in underserved areas addresses a proximity deficit that, in urban settings, may not exist. Stabilizing existing retail in spatially concentrated, low-income-serving locations, through targeted lease support, zoning protections, or community ownership structures, would target the source of this vulnerability. The category heterogeneity identified here sharpens this prescription. Grocery requires location stabilization because the access and substitution channels compound against low-income households, and it is the binding welfare constraint for 55% of neighborhoods. HPC welfare loss is positively associated with the share of households without a vehicle ($\beta_{\text{no_vehicle}} = +0.086$), suggesting that mobile delivery or transportation subsidies could complement store preservation in that category. The threefold variation in income gradient magnitude across metros (-0.61 pp in San Diego to -1.94 pp in San Jose) indicates that income-related welfare disparities differ across metropolitan retail environments, and interventions should be calibrated accordingly. The equilibrium framework developed here can identify which neighborhoods face the greatest exposure under plausible closure scenarios, enabling preemptive targeting rather than reactive response.

These findings connect to several established research trajectories. Recent work has shown that the spatial distribution of human settlements structures mobility patterns[21], and that behavioral mobility data capture dimensions of urban structure that distance-based approaches miss[10,11,13,14,22]. Our equilibrium framework extends these insights from description to counterfactual prediction, showing that consumer reallocation following closures is a structured redistribution that, in grocery, improves welfare through substitution for advantaged neighborhoods and worsens it for disadvantaged ones. This asymmetry arises from the equilibrium structure itself, specifically from how displaced demand is reallocated across surviving alternatives,

and is not visible in static access metrics or reduced-form propagation models. The welfare measurement draws on the structural demand tradition that uses inclusive-value measures to quantify how spatial access shapes urban welfare[5,17], while using equilibrium-share-weighted payoffs to measure the consequences of commercial contraction. Unlike the inclusive-value form, our measure places greater weight on brands that neighborhoods visit more frequently. Our counterfactual design, which separately manipulates spatial environment and demographic characteristics while holding the closure scenario fixed, is a model-based analogue of the place-versus-person distinction in the neighborhood effects literature[23], and in our simulations spatial inputs have a larger influence than demographic inputs.

Several limitations bound these findings. The analysis covers five California metropolitan areas, and the magnitude of the trap may differ in regions with lower density, greater sprawl, or different retail market structure[6]. We estimate a single distance decay coefficient for all neighborhoods. Allowing distance sensitivity to vary with income or vehicle access could change the magnitude of the estimated welfare disparities. Because the equilibrium is resolved at the neighborhood-to-brand level rather than at individual store locations, brand-level congestion may smooth localized crowding effects and understate spatial variation in congestion-related welfare losses. GPS mobility panels exhibit documented demographic coverage biases[24,25]. Using within-category visit shares partially mitigates differences in overall sampling intensity, but does not eliminate brand-specific or spatially varying coverage biases. Noisier visit patterns for lower-income neighborhoods may also affect the estimated income gradient in ways that cannot be determined from the available data. The model also holds category-level visit budgets fixed, excluding substitution across retail categories or shifts toward alternative shopping channels. The model treats retail supply contraction as exogenous. Endogenizing firm entry and exit decisions could reveal whether the trap repels potential entrants through low expected revenue or attracts them through underserved demand, a distinction with direct implications for the persistence of the trap. Despite these limitations, the vulnerability trap defined here appears to be a structural property of how commercial ecosystems are spatially organized relative to the neighborhoods they serve, and the equilibrium framework provides the tools to identify which neighborhoods are most exposed before the next contraction arrives.

Methods

Data sources and geographic scope

We used aggregated monthly visit counts from two data products obtained through the Dewey data marketplace[26]. Monthly mobility patterns come from Advan[27], which processes anonymized GPS signals from approximately 35 million devices across North America to produce origin-destination visit matrices at the census block group (CBG) level. Points of interest (POIs) and their associated attributes come from SafeGraph[28]. The combined dataset spans January 2022 to December 2024 and covers 23 North American Industry Classification System (NAICS)[29] categories. Brands with at least 30 monthly recorded visits whose locations are consistently present across all 36 months were retained, yielding 360 retail brands with sufficient temporal coverage for reliable parameter estimation.

We restricted the analysis to five California metropolitan areas (Los Angeles, San Francisco, San Diego, San Jose, and Sacramento), encompassing 11,252 CBGs and over 4 million possible CBG-brand pairs per month. The equilibrium mechanism at the core of our model, in which consumers substitute across competing brands within each retail category, requires environments where consumers face genuine choices among multiple competitors. Rural and exurban areas, where a single retailer may serve all nearby residents, represent a qualitatively different setting in which substitution options are structurally absent[30]. To evaluate spatial generalization, we trained the model on a stratified subsample of 5,061 CBGs drawn proportionally from each of the five metros (45% per metro, selected randomly within each metro using a fixed seed) and held out the remaining 6,191 CBGs for spatial validation. A separate model trained on all 11,252 CBGs, whose β coefficients agree with the subsample model to three decimal places, was used for all counterfactual simulations (Supplementary Note 4). Mobility data from 2022-2023 were used for training, and 2024 data were reserved as a temporal holdout. Full brand taxonomy, variable definitions, and metro-level summary statistics are reported in Supplementary Note 1.

Brand-level features (location count, visitor volume, mean dwell time, online transaction share) were derived from the mobility and POI data. We also computed an affordability indicator $A_{cb} = 1/(1+|\log y_c - \tau_b^{\text{inc}}|)$, comparing each CBG's log median household income $\log y_c$ to τ_b^{inc}, the visit-weighted mean log income of brand b's observed customers. The indicator is bounded to $[0,1]$, with higher values indicating a closer income match. This indicator is a fixed feature computed once from observed visit data and does not update during equilibrium simulations. CBG-level demographic and socioeconomic variables (median household income, racial composition, vehicle ownership, educational attainment, population density, median housing value) were drawn from the 2024 American Community Survey 5-Year Estimates[31]. GPS mobility panels are known to exhibit demographic coverage disparities and source-dependent distortions[24,25]. Our model is partially robust to this bias because it operates on within-category visit shares rather than absolute counts, so uniform scaling across income groups cancels. Nonetheless, noisier visit patterns for lower-income CBGs may affect the estimated income gradient in ways that cannot be determined from the available data.

Equilibrium model of consumer retail choice

We model retail visitation as a two-level structure following the weak separability assumption standard in discrete choice demand models[32,33]. At the upper level, each CBG c has an observed monthly category budget B_{ck} equal to its total visits to brands in category k, taken as given from the data (see Supplementary Note 6). This assumption is best suited to essential categories, whose demand is inelastic. Consumers do not stop buying groceries when a store closes[34,35]. They redistribute across remaining options. At the lower level, those visits are allocated across brands within the category according to a congestion-adjusted quantal response equilibrium (QRE)[36]. This structure confines substitution to within-category reallocation, isolating the mechanism of interest, namely how consumers redistribute across competing brands of the same type when one brand contracts.

The utility that consumers in CBG c derive from brand b combines an interpretable economic component with a neural residual through a learned convex combination[37,38]

$$U_{cb} = (1-w)\,U_{cb}^{\text{interp}} + w\,f_\theta(\mathbf{z}_{cb}), \tag{1}$$

where $w \in [0,1]$ is a learned mixing weight, f_θ is a feed-forward neural network on feature vector $\mathbf{z}_{cb}$, which concatenates the input features of Equation 2 other than unfamiliarity with a learned 16-dimensional brand embedding (Supplementary Note 4), and the interpretable component is

$$U_{cb}^{\text{interp}} = \alpha_b + \beta_{\text{dist}} \log d_{cb} + \beta_{\text{afford}} A_{cb} + \beta_{\text{size}} \log n_b + \beta_{\text{dwell}} \log \bar{\tau}_b^{\text{dwell}} + \beta_{\text{online}}\, o_b + \beta_{\text{unfam}}\, \Omega_{cb}(\sigma). \tag{2}$$

Here α_b is a brand fixed effect absorbing unobserved brand quality, d_{cb} is the Haversine distance to the nearest POI of brand b, n_b is the number of brand locations, $\bar{\tau}_b^{\text{dwell}}$ is the mean dwell time at brand b, and o_b is its online transaction share, all derived from the mobility and POI data described above. The interpretable component provides economically meaningful, sign-interpretable parameters essential for counterfactual credibility, while the neural residual captures systematic patterns (nonlinear distance effects and interactions among brand and neighborhood features) that a purely parametric specification cannot represent. The learned mixing weight ($w = 0.37$ in the trained model) allows the data to determine the relative contribution of each component.

The last term in Equation 2 captures demographic unfamiliarity when consumers visit retail destinations whose visitor demographics differ from those of their usual shopping environments[10,39–41]. Unlike affordability, the unfamiliarity term Ω_{cb} is endogenous to the equilibrium allocation σ, creating a feedback loop in which demographic sorting reinforces itself as consumers avoid unfamiliar environments. Each brand b induces a visitor demographic profile $\tau_b(\sigma) = \sum_{c'} B_{c'k(b)}\sigma_{c'b}\mathbf{d}_{c'} \big/ \sum_{c'} B_{c'k(b)}\sigma_{c'b}$ that depends on who currently visits it, where $\mathbf{d}_c$ is CBG c's demographic vector (racial composition and log income) and $k(b)$ is brand b's category. Unfamiliarity is then measured as the distance between this visitor profile and the CBG's "demographic comfort zone," defined as the visit-weighted average demographic profile of brands in category k that CBG c historically visited. Specifically, $\Omega_{cb}(\sigma) = \frac{1}{2}\,\text{TV}_{\text{race}}(c,b) + \frac{1}{2}\,D_{\text{income}}(c,b)$, where $\text{TV}_{\text{race}} = \frac{1}{2}\sum_r |\tau_b^r(\sigma) - \text{comfort}_{ck}^r|$ is the total variation distance between the brand's visitor racial composition and the CBG's demographic comfort zone across four racial groups, and $D_{\text{income}} = \frac{2}{3}|\tau_b^{\text{inc}}(\sigma) - \text{comfort}_{ck}^{\text{inc}}|$ is the scaled absolute income distance. Both components are bounded to $[0,1]$ (see Supplementary Note 2). Total variation distance is the natural metric for comparing probability distributions over racial composition, and the equal weighting of the two components is assumption-light because both are already scaled to $[0,1]$. Affordability captures static price-tier matching. Unfamiliarity captures the dynamic demographic comfort of the surviving customer base, which shifts as the equilibrium reorganizes after brand closures.

Given utilities, consumers allocate visits through a category-level quantal response equilibrium[36]

$$\sigma_{cb} = \frac{\exp\left(\lambda_k\left[U_{cb} - \gamma_k \log(1+S_b)\right]\right)}{\sum_{b' \in \mathscr{B}_k} \exp\left(\lambda_k\left[U_{cb'} - \gamma_k \log(1+S_{b'})\right]\right)}, \qquad S_b = \sum_c B_{ck(b)}\,\sigma_{cb}, \tag{3}$$

where $\mathscr{B}_k$ is the set of brands in category k, $\lambda_k > 0$ governs choice sharpness (higher λ corresponds to more deterministic substitution, approaching Nash equilibrium as $\lambda \to \infty$), and $\gamma_k > 0$ is category-specific congestion sensitivity[42,43]. Aggregate patronage S_b is the budget-weighted sum of the allocation and in turn shapes it, producing the fixed-point structure that defines the equilibrium. The log-congestion term is a payoff adjustment rather than a utility shifter, which preserves the interpretation of U_{cb} as a preference index. This congestion structure is what distinguishes the model from a standard multinomial logit. Removing a brand does not simply renormalize shares over remaining options but triggers a cascading reorganization in which displaced consumers crowd surviving brands, altering their attractiveness and prompting further reallocation until the system reaches a new fixed point. Under conditions stated in Supplementary Note 3, the system admits at least one equilibrium allocation, and every reported solve, including each counterfactual, was initialized from uniform within-category shares and re-solved from scratch.

The unfamiliarity term Ω_{cb} and visit shares σ depend on one another. We therefore solve the system with a two-timescale iteration. An outer loop updates unfamiliarity from the current allocation, and an inner loop solves the QRE in Equation 3 with

unfamiliarity treated as fixed. Convergence of this scheme is discussed in Supplementary Note 3. All model parameters, namely the brand fixed effects $\{\alpha_b\}$, the coefficient vector β, the category-level $\{\lambda_k, \gamma_k\}$, the neural parameters θ, and the mixing weight w, were jointly estimated by penalized maximum likelihood on 2022-2023 visit data, holding 2024 out for temporal validation (see Supplementary Note 4 for the loss function, neural network architecture, optimizer specification, and learned parameter values). Because the equilibrium shares are defined implicitly as the fixed point of Equation 3, we approximated parameter gradients by differentiating through a single equilibrium step from the converged fixed point[44,45], without unrolling the fixed-point iterations and keeping training memory independent of iteration count.

Welfare measurement

We measure welfare at the level of realized consumer payoffs rather than raw preferences because commercial contraction reshapes welfare through two mechanisms simultaneously. The first is the direct loss of access when a brand closes. The second is the crowding externality at the surviving alternatives that absorb displaced demand[46,47]. Aggregating at the utility level captures the first mechanism but misses the second because it does not account for how congestion at surviving brands reduces the payoff consumers receive. For each CBG c and brand b we therefore define the congestion-adjusted payoff $V_{cb} = U_{cb} - \gamma_{k(b)} \log(1 + S_b)$, which is the model-implied payoff a consumer in CBG c receives from visiting brand b given the equilibrium crowd S_b at that brand. This is the same payoff that consumers internalize when making choices in Equation 3, so welfare and choice are measured on a single consistent scale.

Experienced welfare at the category level is the aggregation of payoffs across brands, weighted by the CBG's equilibrium visit shares,

$$W_{ck} = \sum_{b \in \mathscr{B}_k} \sigma_{cb} V_{cb}. \tag{4}$$

For any contraction scenario producing a new equilibrium $\{\sigma^{\text{post}}, \mathbf{S}^{\text{post}}\}$, we measure welfare retention as the exponentiated change in experienced welfare,

$$R_{ck} = \exp\left(W_{ck}^{\text{post}} - W_{ck}^{\text{pre}}\right), \tag{5}$$

which equals 1 at baseline. Because R_{ck} depends only on the change in experienced welfare, it is invariant to common additive shifts in utility. A value of $R_{ck} = 0.80$ indicates 80% retention of the exponentiated payoff index relative to that CBG's baseline. Because W_{ck} is share-weighted, changes in payoffs at frequently visited brands have greater influence on welfare retention than changes at rarely visited alternatives. This share-weighted measure is related to the standard inclusive-value welfare form[46,48] through the entropy of the choice distribution. Specifically, $R_{ck} = R_{ck}^{\text{IV}} \cdot \exp(-\Delta H_{ck}/\lambda_k)$, where $R_{ck}^{\text{IV}} = \exp(\text{IV}_{ck}^{\text{post}} - \text{IV}_{ck}^{\text{pre}})$ is the inclusive-value retention and ΔH_{ck} is the change in Shannon entropy of the within-category choice distribution (Supplementary Note 5).

We use experienced welfare rather than the inclusive-value form to focus on payoffs weighted by equilibrium visitation, rather than the option value of the entire choice set. Share-weighting concentrates the measure on the brands consumers visit most, and the congestion adjustment captures the welfare reduction at surviving brands from the crowds that closures redirect there (Supplementary Note 5). This share-weighted measure is not necessarily monotone in the available choice set, so positive substitution contributions should not be interpreted as increases in conventional consumer surplus.

The four essential retail categories in our analysis, collectively denoted $\mathscr{K}$ (grocery, HPC, gasoline stations, and general merchandise), differ in both their functional role for households and their vulnerability to contraction. Adequate grocery access does not compensate a neighborhood that loses most of its pharmacy options[49,50]. We therefore use the weakest-link aggregation $R_{\min,c} = \min_{k \in \mathscr{K}} R_{ck}$ as the primary measure of essential retail vulnerability per neighborhood. The dependent variable in the pooled and per-metro regressions is the complementary loss index $1 - R_{\min,c}$, representing the largest modeled loss in welfare retention across essential categories, and the per-category regressions use $1 - R_{ck}$.

Counterfactual simulations

We applied three distinct simulation designs across the four essential retail categories using the model trained on all 11,252 CBGs. Every scenario re-solves the complete two-timescale system with no endogenous component held fixed, so counterfactual predictions reflect the full structural response of consumer choice to supply shocks.

Category-proportional contraction. We progressively removed the lowest-traffic POIs within each essential category at five contraction levels ($X \in \{10\%, 20\%, 30\%, 40\%, 50\%\}$). POI traffic was ranked using observed monthly visit counts, and the bottom $X\%$ of POIs within each category were closed simultaneously. This design ensures each essential service type experiences the same intensity of contraction, reflecting a pattern where retail closures are driven by category-specific market

forces. Closing a POI of brand b affects each CBG c through a rescaling of its accessibility to brand b, based on a Hansen-style accessibility index[51]

$$\Delta U_{cb}^{\text{access}} = |\beta_{\text{dist}}| \cdot \log\left(\frac{a_{cb}^{\text{post}}}{a_{cb}^{\text{pre}}}\right), \qquad a_{cb} = \sum_{p \in \mathscr{P}_b} \frac{1}{d_{cp} + \varepsilon}, \tag{6}$$

where $\mathscr{P}_b$ is the set of open POIs of brand b, d_{cp} is the Haversine distance from CBG c to POI p in km, and $\varepsilon = 0.1$ km prevents division by zero. If all POIs of a brand close, the brand is removed from the choice set entirely. The accessibility index generalizes beyond nearest-POI distance by accounting for all of a brand's locations, and the log-ratio formulation uses the estimated distance decay parameter to scale the utility adjustment. For each contraction level, the modified utilities enter the full two-timescale equilibrium solver, which endogenously recomputes unfamiliarity and aggregate patronage to produce the post-contraction allocation.

Input-level counterfactual swaps. To compare the influence of spatial exposure and demographic composition on the modeled welfare gap between disadvantaged and advantaged neighborhoods[5,23], we replaced model inputs of CBGs in the lowest within-metro income quintile with the corresponding within-metro median values of CBGs in the highest quintile and re-solved the full equilibrium from baseline through all five contraction levels. Two swaps were performed. CF-A replaces each lowest-income CBG's distance vector $\{d_{cb}\}_b$ and the corresponding accessibility scaling with the within-metro median values of highest-income CBGs. CF-B replaces each lowest-income CBG's income, racial composition, and demographic comfort-zone parameters with the within-metro median values of highest-income CBGs, which flow through both the affordability and unfamiliarity terms. Swaps operate on model inputs rather than utility components, so they propagate through both the interpretable and neural components of Equation 1, testing the full structural response rather than an isolated coefficient. Gap closure is measured as the fraction of the $R_{\min,c}$ welfare gap between lowest- and highest-income CBGs eliminated by each swap.

Brand-preserving contraction with ordered closure. To test whether retail contraction exhibits the market-selection, gap-widening, and recovery-asymmetry patterns characteristic of the vulnerability trap[19,20], we designed a contraction scenario that preserves brand-level presence. Each brand b loses 50% of its POIs over ten progressive steps (5% per step). No brand disappears entirely, isolating the spatial-redistribution mechanism from brand-composition effects. Within each brand, three closure orderings were applied. *Traffic-based* closes the lowest-traffic POIs first, approximating market-driven exit. *Competition-based* closes POIs with the highest count of same-category competitors within 10 km first, approximating strategic rationalization of redundant locations. *Random* closes POIs in uniformly random order, averaged over five realizations, and serves as a null benchmark. All three orderings reach the same final contraction level but close different sets of POIs along different trajectories, and contrasts across orderings show how different closure-selection rules affect the resulting welfare disparities. After the closure phase, each simulation continued with a ten-step recovery phase in which POIs reopen in reverse order of closure. Under traffic-based ordering, this approximates market-driven recovery in which the highest-traffic locations reopen first. We measure recovery asymmetry by comparing welfare regained by neighborhoods in the lowest and highest income quintiles when half of the closed POIs have reopened.

Data Availability

For sociodemographic information at the census block group level, the 2024 5-year American Community Survey (ACS) was used (American Community Survey (ACS)). All the census block group and metropolitan shapefiles come from TIGER/Line datasets (TIGER/Line Shapefiles). The places and mobility data used in our study were sourced from SafeGraph (SafeGraph) and Advan (Advan Research Corporation) through a subscription to the Dewey data marketplace (Dewey), and are therefore not publicly available. Additional details regarding the datasets and subscription services are available at deweydata.io.

Code Availability

The analysis was conducted using custom code in Python (version 3.13). The model source code, training pipeline, and simulation scripts are available on GitHub. For further inquiries, please contact the corresponding author.

References

1. 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).
2. Lopez, R. A., Marchesi, K. & Steinbach, S. Dollar store expansion and independent grocery retailer contraction. *Appl. Econ. Perspectives Policy* **46**, 514–533 (2024).

3. 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).
4. Mathis, W. S. *et al.* Vulnerability index approach to identify pharmacy deserts and keystone pharmacies. *JAMA Netw. Open* **8**, e250715, DOI: 10.1001/jamanetworkopen.2025.0715 (2025).
5. Allcott, H. *et al.* Food deserts and the causes of nutritional inequality. *Q. J. Econ.* **134**, 1793–1844, DOI: 10.1093/qje/qjz015 (2019).
6. Schuetz, J., Kolko, J. & Meltzer, R. Are poor neighborhoods "retail deserts"? *Reg. Sci. Urban Econ.* **42**, 269–285 (2012).
7. Qato, D. M. *et al.* 'pharmacy deserts' are prevalent in chicago's predominantly minority communities, raising medication access concerns. *Heal. Aff.* **33**, 1958–1965 (2014).
8. Handbury, J. Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. *Econometrica* **89**, 2679–2715, DOI: 10.3982/ECTA11738 (2021).
9. Bower, K. M., Thorpe, R. J., Rohde, C. & Gaskin, D. J. The intersection of neighborhood racial segregation, poverty, and urbanicity and its impact on food store availability in the United States. *Prev. Medicine* **58**, 33–39, DOI: 10.1016/j.ypmed.2013.10.010 (2014).
10. Moro, E., Calacci, D., Dong, X. & Pentland, A. Mobility patterns are associated with experienced income segregation in large us cities. *Nat. communications* **12**, 4633 (2021).
11. García Bulle Bueno, B. *et al.* Effect of mobile food environments on fast food visits. *Nat. Commun.* **15**, 2291 (2024).
12. Abbiasov, T. *et al.* The 15-minute city quantified using human mobility data. *Nat. Hum. Behav.* **8**, 445–455 (2024).
13. Xu, F. *et al.* Using human mobility data to quantify experienced urban inequalities. *Nat. human behaviour* **9**, 654–664 (2025).
14. Yabe, T., García Bulle Bueno, B., Frank, M. R., Pentland, A. & Moro, E. Behaviour-based dependency networks between places shape urban economic resilience. *Nat. human behaviour* **9**, 496–506 (2025).
15. Chang, S. *et al.* Mobility network models of covid-19 explain inequities and inform reopening. *Nature* **589**, 82–87 (2021).
16. Berry, S., Levinsohn, J. & Pakes, A. Automobile prices in market equilibrium. *Econometrica* **63**, 841–890, DOI: 10.2307/2171802 (1995).
17. Miyauchi, Y., Nakajima, K. & Redding, S. J. The economics of spatial mobility: Theory and evidence using smartphone data. *The Q. J. Econ.* **140**, 2507–2570 (2025).
18. Rite Aid Corporation. Rite Aid Takes Steps to Accelerate Transformation and Position Company for Long-Term Success. Business Wire (2023).
19. Azariadis, C. & Stachurski, J. Poverty traps. *Handb. economic growth* **1**, 295–384 (2005).
20. Kraay, A. & McKenzie, D. Do poverty traps exist? Assessing the evidence. *J. Econ. Perspectives* **28**, 127–148, DOI: 10.1257/jep.28.3.127 (2014).
21. Boucherie, L., Maier, B. F. & Lehmann, S. Decoupling geographical constraints from human mobility. *Nat. Hum. Behav.* **9**, 2564–2575 (2025).
22. Gazmeh, H., Guo, Y. & Qian, X. Understanding the opportunity-centric accessibility for public charging infrastructure. *Transp. Res. Part D: Transp. Environ.* **131**, 104222 (2024).
23. Chetty, R. & Hendren, N. The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects. *The Q. J. Econ.* **133**, 1107–1162, DOI: 10.1093/qje/qjy007 (2018).
24. Gallotti, R., Maniscalco, D., Barthelemy, M. & De Domenico, M. Distorted insights from human mobility data. *Commun. Phys.* **7**, 421 (2024).
25. Coston, A. *et al.* Leveraging administrative data for bias audits: Assessing disparate coverage with mobility data for COVID-19 policy. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, 173–184, DOI: 10.1145/3442188.3445881 (ACM, 2021).
26. Dewey Data. Dewey | Academic Research Data. https://deweydata.io (2024).
27. Advan Research. Foot traffic / monthly patterns [dataset]. https://doi.org/10.82551/BEB1-2831 (2022).
28. SafeGraph. Global places (poi) & geometry [dataset]. https://doi.org/10.82551/SMXB-1K04 (2022).
29. US Census Bureau. North American Industry Classification System (NAICS). https://census.gov/naics (2022).

30. Clarke, I. *et al.* Retail competition and consumer choice: contextualising the "food deserts" debate. *Int. J. Retail. & Distribution Manag.* **32**, 89–99 (2004).

31. US Census Bureau. American Community Survey (ACS). https://census.gov/programs-surveys/acs (2024).

32. 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).

33. Berry, S. T. Estimating discrete-choice models of product differentiation. *RAND J. Econ.* **25**, 242–262 (1994).

34. Kotschedoff, M. J., Kowalczyk, L. & Breugelmans, E. The persistence of grocery shopping behavior and retailer choice: Evidence from a major labor strike. *Quant. Mark. Econ.* **23**, 447–490 (2025).

35. Shi, Q., Inman, J. J. & Gauri, D. "store closed": How consumers respond to retail store closures. *Mark. Sci. Inst. Work. Pap. Ser.* (2023).

36. McKelvey, R. D. & Palfrey, T. R. Quantal response equilibria for normal form games. *Games economic behavior* **10**, 6–38 (1995).

37. Hartford, J., Lewis, G., Leyton-Brown, K. & Taddy, M. Deep IV: A flexible approach for counterfactual prediction. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, vol. 70 of *Proceedings of Machine Learning Research*, 1414–1423 (PMLR, 2017).

38. Farrell, M. H., Liang, T. & Misra, S. Deep neural networks for estimation and inference. *Econometrica* **89**, 181–213, DOI: 10.3982/ECTA16901 (2021).

39. Davis, D. R., Dingel, J. I., Monras, J. & Morales, E. How segregated is urban consumption? *J. Polit. Econ.* **127**, 1684–1738, DOI: 10.1086/701680 (2019).

40. de la Prada, À. G. & Small, M. L. How people are exposed to neighborhoods racially different from their own. *Proc. Natl. Acad. Sci.* **121**, e2401661121 (2024).

41. Massenkoff, M. & Wilmers, N. Rubbing shoulders: Class segregation in daily activities. *J. Public Econ.* **244**, 105335 (2025).

42. Daganzo, C. F. & Sheffi, Y. On stochastic models of traffic assignment. *Transp. Sci.* **11**, 253–274, DOI: 10.1287/trsc.11.3.253 (1977).

43. Anderson, S. P., de Palma, A. & Thisse, J.-F. *Discrete Choice Theory of Product Differentiation* (MIT Press, Cambridge, MA, 1992).

44. Bai, S., Kolter, J. Z. & Koltun, V. Deep equilibrium models. *Adv. neural information processing systems* **32** (2019).

45. Blondel, M. *et al.* Efficient and modular implicit differentiation. *Adv. neural information processing systems* **35**, 5230–5242 (2022).

46. Small, K. A. & Rosen, H. S. Applied welfare economics with discrete choice models. *Econom. J. Econom. Soc.* 105–130 (1981).

47. 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).

48. Train, K. E. *Discrete Choice Methods with Simulation* (Cambridge University Press, Cambridge, UK, 2009), 2 edn.

49. Sen, A. *Development as Freedom* (Alfred A. Knopf / Oxford University Press, New York / Oxford, 1999).

50. Alkire, S. & Foster, J. Counting and multidimensional poverty measurement. *J. Public Econ.* **95**, 476–487, DOI: 10.1016/j.jpubeco.2010.11.006 (2011).

51. Hansen, W. G. How accessibility shapes land use. *J. Am. Inst. Planners* **25**, 73–76 (1959).

Funding

Not applicable.

Author information

Authors and Affiliations

Department of Civil and Environmental Engineering, Rice University, Houston, TX

Hossein Gazmeh, Xinwu Qian

Contributions
H.G. and X.Q. conceptualized the study. H.G. curated the data. H.G. and X.Q. conducted the formal analysis and developed the methodology. H.G. and X.Q. wrote the manuscript. H.G. and X.Q. contributed to reviewing and editing the final version.

Corresponding author
Correspondence to Xinwu Qian.

Ethics declarations

Ethics approval and consent to participate
Not applicable. This study did not involve human participants, human material, or animals. The analysis relies on aggregated, de-identified human mobility data obtained commercially from SafeGraph and Advan through the Dewey data marketplace, together with publicly available aggregate statistics from the U.S. Census Bureau (American Community Survey and TIGER/Line). No individual-level or identifiable data were collected or analyzed, and no intervention involving human subjects was performed. Ethics approval and informed consent to participate were therefore not required.

Consent for publication
Not applicable. The manuscript does not contain any data, images, or details relating to an individual person.

Competing interests
The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- NatCommCommercialResilienceshared.pdf
- NatCommSI.pdf