Endogenous Amenities and Endogenous Markets: Evidence from Baltimore

Samsun Knight* Jordan Hutchings†

University of Toronto University of Toronto

August 18, 2026

Abstract

Restaurants attract residents, and residents attract restaurants. This two-sided sorting confounds observational estimates of consumer preferences and may lead standard restaurant entry models to systematically mispredict which areas are best for potential new entrants. We estimate a joint model of restaurant entry and household location choice in the Baltimore metropolitan area, exploiting the 2024 Francis Scott Key bridge collapse and Baltimore County's four-year zoning cycle as quasi-experiments to identify consumer demand and household sorting preferences. The bridge collapse reveals that consumers are meaningfully sensitive to travel time, with demand decreasing by 3.9% for every 10% increase in travel time, while the zoning event study identifies a non-monotonic preference for living near restaurants, with households valuing restaurant access at moderate distance but disliking living too close. We then embed these quasi-experimental estimates into a dynamic entry model and show that an endogenous-market specification, which anticipates household sorting in response to area amenities, significantly reduces out-of-sample prediction error relative to a standard exogenous-dynamics baseline. Counterfactual analysis of transit improvements further shows that the two specifications rank candidate road improvements differently, and project different long-run profitability effects for nearby business locations. Taken together, our results show that "crowded" markets are often not as saturated as they appear, as higher competitive density can make neighborhoods more attractive places to live and expand local demand in the long-term.

Keywords: restaurant entry, household sorting, spatial demand, transportation infrastructure, zoning

*samsun.knight@rotman.utoronto.ca

†jordan.hutchings@rotman.utoronto.ca

1 Introduction

Where should a restaurant locate? The conventional answer relies on estimates of market size that are exogenous to entry, where area population dynamics are the same regardless of which firms enter (Bresnahan and Reiss 1991, 1994; Seim 2006; Aguirregabiria and Mira 2007; Bajari, Benkard, and Levin 2007). But in practice, households move toward neighborhoods with appealing amenities, including restaurants (Diamond 2016; Cao et al. 2024). If restaurants anticipate this feedback, then a thicker local competitor environment may provide an offsetting benefit to other potential entrants, as a greater density of area restaurants may carve up today's local market into more slices, but expand the overall pie over the longer term. In the limit, this feedback may even turn competition into a net complementarity, as a cluster of restaurants could collectively draw enough new residents nearby that they all benefit from each other's presence.

Modeling this feedback requires solving an identification problem that standard entry frameworks sidestep by treating population dynamics as exogenous. When restaurants choose locations and households choose neighborhoods endogenously, observational estimates of consumer and household preferences may be confounded by restaurants locating strategically relative to their customer base, and estimated models of local restaurant demand may be analogously biased by household sorting over time (Dingel and Tintelnot 2026). But if this estimation challenge can be surmounted, credibly modeling this full feedback loop may offer a potent new opportunity to improve on standard practices, both for businesses choosing where to open locations and for policymakers designing spatial economic development policy.

In this paper, we estimate a joint model of restaurant entry and household location choice in the Baltimore metropolitan area, exploiting the 2024 collapse of the Francis Scott Key bridge and Baltimore County's four-year zoning change cycle as quasi-experiments to obtain credible estimates of consumer demand and household sorting preferences. We then show that an "endogenous-market" restaurant entry model, in which restaurants anticipate how households move towards amenities, significantly outperforms standard benchmark models, particularly in predicting which restaurants will survive.

First, to overcome the above-mentioned biases from two-sided spatial sorting of restaurants and households, we use two quasi-experiments to establish credible estimates of how travel distance affects contemporaneous restaurant demand and how local restaurant density affects household location choice. We estimate consumer sensitivity to travel distance using the unexpected collapse of the Francis Scott Key bridge near Baltimore in 2024 as an exogenous shock to area travel time,[1] showing in reduced-form difference-in-differences analysis that restaurants nearest to the bridge saw a sharp monthly spending decline of roughly 4% after the collapse. This decline started directly after the event and sustained over the following 9 months; furthermore, we find similar or larger effects in areas moderately far from the bridge that primarily suffered negative traffic-congestion spillovers, supporting our interpretation of the event as a travel time shock. We use indirect inference (Gourieroux, Monfort, and Renault 1993; Gechter and Tsivanidis 2025; Andrews, Gentzkow, and Shapiro 2020) to leverage this quasi-experimental variation to identify preferences over travel time in a structural model of restaurant demand, recovering an elasticity of demand to travel time of -0.39. We then estimate household preferences to live near restaurants using Baltimore County's four-year zoning cycle, which in 2016 reclassified a large number of parcels into restaurant-permitting zones on a fixed statutory schedule and created sharp, exogenously-timed variation in restaurant density.[2] We show that this policy shock was preceded by null pre-trends in both restaurant count and population levels, and then afterwards 1) rezoned areas saw significantly increased local restaurant entry, 2) areas with increased immediate restaurant density (within 1 mile) saw significant population declines, and 3) areas with increased density

1. On March 26, 2024, a container ship struck the bridge, suddenly increasing driving times in the area by removing a central connection in the area's transit network, both through direct route changes and through congestion spillovers to other adjacent routes (Chang et al. 2025).

2. Parcel selection into rezoning is not random, but the four-year timing of rezoning is set exogenously by policy, and as such, any persistent confounds would be expected to show up as significant pre-trends; the absence of such pre-trends thus supports the validity of our estimates. We also show that non-restaurant businesses had null changes, suggesting that these are restaurant-specific amenity responses.

at intermediate distance (1 to 2 miles) saw significant population increases, consistent with households disliking living very close to restaurants (noise and congestion disamenity) but liking living moderately close (better restaurant access). Then, as above, we leverage these reduced-form effects to identify preferences for living near restaurants in a structural model of household location choice using indirect inference.

With these preference estimates in hand, we develop a dynamic restaurant entry model that explicitly incorporates these endogenous household-sorting dynamics, both in order to quantify the importance of these dynamics for researchers and to clarify takeaways for businesses and policymakers. Specifically, we directly embed the above consumer-demand and household-location models into an entry model based on Aguirregabiria and Mira (2007), allowing firms to anticipate how restaurant density will affect population longer-term, and then compare this model's performance against a benchmark exogenous-market entry model where expectations of population dynamics are set from local area projections from the Maryland Department of Planning. The endogenous-market specification reduces total out-of-sample MSE by 5–8% relative to this exogenous-market benchmark, an improvement driven specifically by the restaurant→population feedback channel. The two specifications also imply substantively different effects of candidate road improvements; by overlooking the endogenous feedback-loop dynamics of urban growth, the exogenous specification overprioritizes areas that look underserved by current restaurant counts and overlooks areas where transit interventions could kick off a virtuous cycle of drawn-in businesses and drawn-in population. This misprioritization can be severe enough to reverse the ranking of candidate corridors: a policymaker choosing which single corridor to build would select a different road depending on which specification's projections they used. Analogously, this analysis implies that firms attempting to use exogenous-population models to assess profitability in a changing urban landscape may systematically mis-measure how large-scale neighborhood changes, such as major new roads, will affect the long-run profitability of affected locations.

Taken together, our results provide clear evidence that restaurant entry and residential location are jointly determined, and that this implies a novel complementarity effect from higher local competitor density, through drawing more households to the area over the longer-term, that previous models have overlooked. For restaurant managers, we show that site selection that incorporates these endogenous neighborhood dynamics can meaningfully improve predictions of which areas are most likely to support restaurants long-term, as "crowded" markets may not always be as saturated as they seem. For policymakers, our counterfactual analysis of transit improvements shows that standard models can misestimate entry responses and rank candidate roads differently once endogenous dynamics are accounted for. For researchers, our results demonstrate that naive observational estimates of preferences that are entangled with spatial sorting can be severely biased, but also that models that credibly account for the endogenous evolution of urban areas may provide a material improvement over current standard approaches.

The remainder of the paper proceeds as follows. Section 2 reviews related literature. Section 3 describes the data. Section 4 presents the consumer demand model. Section 5 develops the household location choice model. Section 6 estimates the restaurant entry model, compares specifications, and presents a counterfactual analysis of transit improvements. Section 7 concludes.

2 Related Literature

This paper contributes to three literatures, namely the urban economics literature on endogenous amenities and consumption geography, the empirical IO and marketing literature on dynamic firm entry, and the broader public economics literature on the effects of zoning and infrastructure on local economic activity.

First, our work builds on the urban economics literature on endogenous amenity dynamics. This literature includes Diamond (2016), who estimates how skill sorting across cities affects local amenity provision and welfare, and Couture et al. (2024), who show that high-income households' concentration in urban cores generates amenity improvements that reinforce spatial sorting; as well as consumption geography studies about how consumers value spatially distributed amenities, such as Miyauchi, Nakajima, and Redding (2025), Davis et al. (2019), Fox, Postrel, and McLaughlin (2007), and Cao et al. (2024), which generally rely on either observational relationships or spatial instruments like distance-to-central-business-district in order to infer preferences. Most closely related to the present work, Almagro and Domínguez-Iino (2025) jointly estimate household location choice and retail amenities at the city-neighborhood level in Amsterdam, demonstrating that amenity endogeneity substantially affects welfare calculations from urban policies. Broadly speaking, this literature develops macro- or city-scale evidence that amenities are endogenous to household entry and also then induce further entry, with more restaurants, shops, or services making a location unambiguously more attractive. In our much more granular analysis, we use a bridge collapse and sharp shift in zoning policy to reveal a more nuanced pattern at a highly local level, with restaurants serving as amenities at moderate distance but as disamenities at immediate proximity, possibly reflecting disutility from noise and congestion,[3] and integrate these dynamics into a restaurant entry framework to develop implications for business location decisions.

Related to this, there is a long empirical IO literature on dynamic entry, both for retail more broadly and for restaurants in particular, that has developed computationally tractable methods for estimating models of firm entry and exit decisions. Among many others, this includes Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Seim (2006), and Singh, Hansen, and Blattberg (2006), Igami and Yang (2016), Blevins, Khwaja, and Yang (2018), Nishida (2015), Arcidiacono et al. (2016), and Caoui, Hollenbeck, and Osborne (2025). To the best of our knowledge, this literature, while it richly accounts for strategic dynamics, to date has always treated market size dynamics as exogenous; our

3. Observational hedonic evidence of a similar non-monotonic distance–value pattern for commercial activity has been documented by Matthews and Turnbull (2007) and Yang, Song, and Choi (2016); to the best of our knowledge, ours is the first quasi-experimental evidence of this pattern.

contribution is to join the above endogenous amenities literature to this dynamic entry IO literature, integrating quasi-experimentally identified consumer demand and household location choice models into a dynamic entry framework estimated via nested pseudo-likelihood (Aguirregabiria and Mira 2007), where long-term local market size is itself endogenous to area amenity density. We demonstrate that incorporating this endogeneity improves out-of-sample prediction of restaurant entry and exit and produces meaningfully different counterfactual policy predictions.

Third, the present work relates to prior studies using zoning and infrastructure shocks to study the effect of local urban policy on economic activity. For example, Nishida (2014), Datta and Sudhir (2013), and Suzuki (2013) exploit land-use regulation to study the effects of zoning on firm entry, while Donaldson and Hornbeck (2016), Allen and Arkolakis (2022), and Heblich, Redding, and Sturm (2020) use transportation investments to study how transit determines economic geography. We exploit an unexpected major bridge collapse and a sharp zoning policy shock to provide new evidence on both transit and zoning effects, revealing that restaurants are disamenities at immediate proximity but amenities at moderate distance, and that consumers have significant disutility from travel time that is substantially misestimated without credible quasi-experimental variation.[4]

3 Data

For our full suite of analyses, we assemble five data sources covering consumer spending, restaurant entry and exit, driving distances, land-use regulation, and demographics in the Baltimore–Washington metropolitan area. This section summarizes each source; for further details and full construction procedures, see Appendix A.

4. Our event studies of the bridge collapse and zoning shock also contribute to a small literature assessing the predictive power of quantitative spatial models by comparing model-predicted changes to observed outcomes using plausibly exogenous shocks (Ahlfeldt et al. 2015; Monte, Redding, and Rossi-Hansberg 2018; Adão, Arkolakis, and Esposito 2026; Dingel and Tintelnot 2026); in line with that literature, we show that observational estimates of elasticities and other economically relevant parameters that are entangled with spatial sorting are highly likely to be biased.

Our primary source of consumer demand data is the SafeGraph Spend panel (SafeGraph 2022). SafeGraph Spend links a large store location panel with card transaction records from a major US transaction data aggregator to produce establishment-level monthly aggregates of spending, transaction counts, and customer counts. We restrict to restaurant establishments (NAICS codes starting with 722) located within 60 minutes driving time of the Francis Scott Key Bridge. This gives us a balanced panel of 2,629 restaurants from January 2023 through December 2024, spanning the period before and after the bridge collapse on March 26, 2024.[5] Panel-level summary statistics are reported in Appendix Table A1.

Second, we construct a restaurant entry-and-exit panel from Data Axle's Business Academic historical database. Data Axle compiles its records from business registrations, telephone directories, and proprietary verification processes, and tracks establishments longitudinally via persistent identifiers. The database is widely used in research on retail dynamics (e.g., Cao et al. 2024) and provides comprehensive coverage of the restaurant sector. We use annual snapshots from 2012 through 2020 to measure entry and exit from 2013 to 2019 for restaurant establishments (as above, defined by NAICS code) in a 10-county study area centered on Baltimore,[6] aggregated to entry and exit counts at the Census block group (CBG) level. We observe 16,440 entries and 13,930 exits across 3,106 CBGs from 2013 to 2019. Entry/exit panel summary statistics are reported in Appendix Table A2.

We note that the consumer demand model uses 2023–2024 spending data (to exploit the 2024 bridge collapse), while the household sorting and entry models use 2013–2019 data (to exploit the 2016 zoning change) due to data availability constraints across different providers and time periods. When combining these estimates in our final analyses, we are implicitly

5. We note that the SafeGraph panel is predominantly based on debit card transactions, and as debit card use is more common among lower-income populations, revenue levels may tend to be lower in this panel than would be found in a pure random sample; however, since we include extensive fixed effects in all specifications, including restaurant and time fixed effects, any such level bias would be differenced out of our effect estimates.

6. The study area comprises the seven counties of the official Baltimore-Columbia-Towson CBSA (Anne Arundel, Baltimore City, Baltimore County, Carroll, Harford, Howard, and Queen Anne's) plus three adjacent counties (Kent, Montgomery, and Prince George's). We include the adjacent counties to provide additional control locations for the household sorting model and because they contain a meaningful share of the restaurants in the bridge collapse demand sample.

assuming that the relevant preferences are stable across these periods. In Appendix C.3, we provide observational evidence supporting this assumption, showing that observational household preference ratios are essentially identical when the sorting model is re-estimated on 2022–2023 ACS data, and that the observational relationship between travel distance and spending is highly qualitatively similar in both 2019 and 2023-2024.

Third, we gather driving distances and durations data from OSRM (Open Source Routing Machine, Luxen and Vetter (2011)), an open-source routing engine that computes shortest paths on OpenStreetMap road network data. We obtain both pre-collapse and post-collapse driving distances by computing distances and durations both before and after removing the Francis Scott Key bridge connection from the transit network. We also augment these route-based driving time changes with congestion effect estimates from Chang et al. (2025), who offer an unusually well-suited complement to our analysis, in that they estimate transit congestion effects from this precise Francis Scott Key bridge collapse across different areas of the Baltimore region; we integrate their estimates with our driving network data to compute the implied congestion-adjusted driving time change from the bridge collapse for each restaurant–home pair. For more details, see Appendix A.2.1.

Fourth, we use digitized parcel-level zoning maps from the 2012 and 2016 cycles of the Baltimore County Comprehensive Zoning Map Process (CZMP). We focus on Baltimore County for this exercise because this county conducts a comprehensive rezoning on a fixed four-year statutory cycle, during which all parcels in the county are subject to potential reclassification, and furthermore because they make their complete zoning maps publicly available for both the 2012 and 2016 cycles. We identify 386 parcels newly classified into restaurant-permitting zones between these 2012 and 2016 maps, and use these changes to compute, for each of 502 CBGs in Baltimore County, the change in restaurant-zoned land area within 1-, 2-, and 3-mile annuli, as well as the change in restaurant-zoned land for each CBG.

Fifth, we obtain CBG-level demographic characteristics from the American Community

Survey (ACS) 5-year estimates, published annually by the US Census Bureau. The ACS pools survey responses over rolling five-year windows to produce reliable small-area estimates at the block group level; we use these to measure CBG population, as well as Census covariates including median household income and population density, and also county-to-county migration flows that inform the household sorting model. Summary statistics are reported in Appendix Table A3.

4 Consumer Demand

The principal goal of this paper is to develop a comprehensive model of restaurant entry with endogenous markets. The first ingredient of that model is a credible model of contemporaneous consumer demand, to capture how demand will vary across potential candidate locations in any given period. However, as noted above, the two-sided spatial sorting of both restaurants and consumers implies that observational variation in demand will be confounded, since restaurants rationally locate where demand is highest, potentially leading to spurious correlations between average travel distance and demand.

We therefore begin our analysis by rigorously establishing the effect of travel time on restaurant consumer demand using a quasi-experiment, and then leverage these quasi-experimental estimates to identify a structural consumer demand model.

4.1 Identification: The Francis Scott Key Bridge Collapse

On March 26, 2024, the container ship *Dali* lost power while departing the Port of Baltimore and struck the Francis Scott Key Bridge, causing its complete collapse. The bridge carried Interstate 695 across the Patapsco River and served approximately 30,000 vehicles daily (Frittelli, Goldman, and Lohman 2024). Its sudden collapse created sharp and substantial exogenous variation in travel times between consumers and restaurants throughout the Baltimore metropolitan area, both through the direct change to the transit network and

Figure 1 — Francis Scott Key Bridge and Study Area

(A) Bridge Location **(B) Study Area**

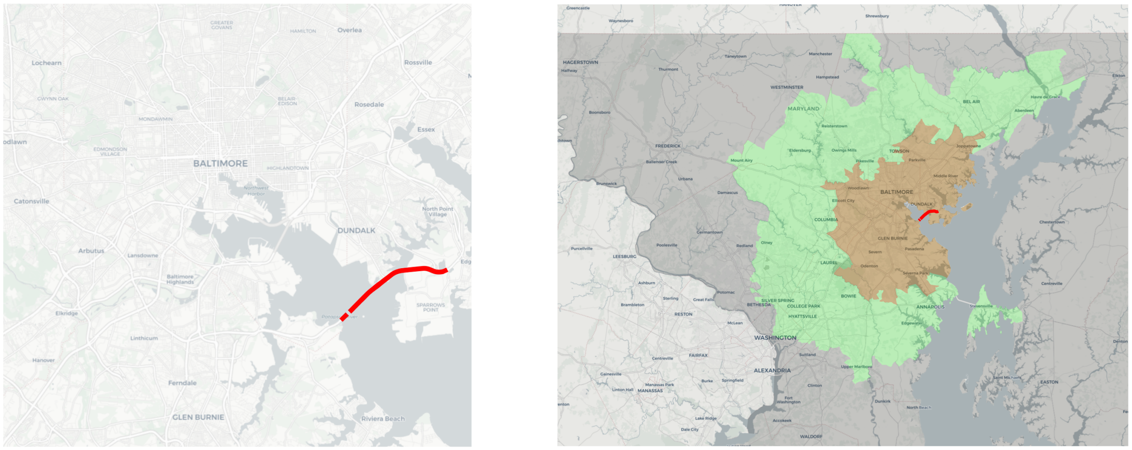

Notes. Panel (A) shows the location of the Francis Scott Key Bridge (red) crossing the Patapsco River south of Baltimore. Panel (B) shows the study area used for the consumer demand model. The light brown region indicates restaurants within 0–30 minutes driving time of the bridge and the green region indicates restaurants 30–60 minutes away. Note that Panel (B) uses county-level shading as a visual approximation; actual treatment assignment is based on each restaurant's individual driving time to the bridge.

through negative congestion spillovers (Chang et al. 2025). Figure 1 shows the bridge location and our study area.

We use differences-in-differences estimation to infer the causal impact of the bridge collapse on restaurant demand. Specifically, to measure the overall impact, we estimate the following equation:

$$\text{Spend}_{jt} = \gamma_j + \delta_t + \beta^{\text{DiD}} \cdot \mathbf{1}_{\text{Treat}_j \times \text{Post}_t} + \varepsilon_{jt} \tag{1}$$

where Spend_{jt} is monthly spending at restaurant j in month t, $\mathbf{1}_{\text{Treat}_j \times \text{Post}_t}$ is an indicator for treatment post-bridge collapse, γ_j are restaurant fixed effects absorbing time-invariant restaurant characteristics, and δ_t are time (month) fixed effects absorbing aggregate shocks to restaurant demand. For the event studies, we estimate an analogous specification with dummies for treatment status interacted with relative-time dummies, to capture both the pre-trend and post-period effects.

In this reduced-form analysis, we define treatment using a simple heuristic, as restaurants within 30 minutes of driving time from the bridge, an approximate threshold consistent with the spatial extent of congestion effects documented in Chang et al. (2025) (see Section A.2.1). While imperfect, this descriptive definition allows us to capture the average effect of the bridge collapse on the restaurants we expect to be the most impacted by the event with minimal assumptions. We define control restaurants as either all restaurants 30–60 minutes of driving time from the bridge, or as all restaurants 45–60 minutes of driving time from the bridge. To the extent that restaurants beyond 30 or 45 minutes of driving time from the bridge may also be impacted by the bridge collapse, the estimated treatment impact in this specification may be attenuated; however, this only further motivates our later use of a structural model to infer the travel time disutility from this quasi-experimental evidence, as there our estimated restaurant demand model can directly model the treatment impact for all restaurants in the sample, including partial treatment for "control" restaurants, when matching the reduced-form results.

Results from this event study are presented in Figure 2. In Panel (A), we present the results using the full sample of restaurants, where the control group is defined as all restaurants 30–60 minutes of driving time from the bridge, and in Panel (B), we present the results using the restricted sample of restaurants, where the control group is defined as all restaurants beyond 45 minutes of driving time from the bridge. In both cases, we recover highly similar results, with flat and statistically null pre-trends followed by a sharp, statistically significant decline in restaurant revenue after the bridge collapse for treated restaurants. Effect sizes correspond to a roughly 4% decline in restaurant revenue compared to the pre-period average and sustain throughout the post-period. We remark that the point estimates are modestly larger when using the farther control group, consistent with a minor attenuation from partial treatment of restaurants 30-45 minutes from the Francis Scott Key Bridge (and also feature slightly larger confidence bands, likely arising from the smaller sample size of restaurants).

Aggregate effect estimates, across a set of alternate window lengths and control group definitions, are reported in Table 1. Results are consistent across control group definitions

Figure 2 — Effect of Bridge Collapse on Restaurant Spending

(A) Full Sample

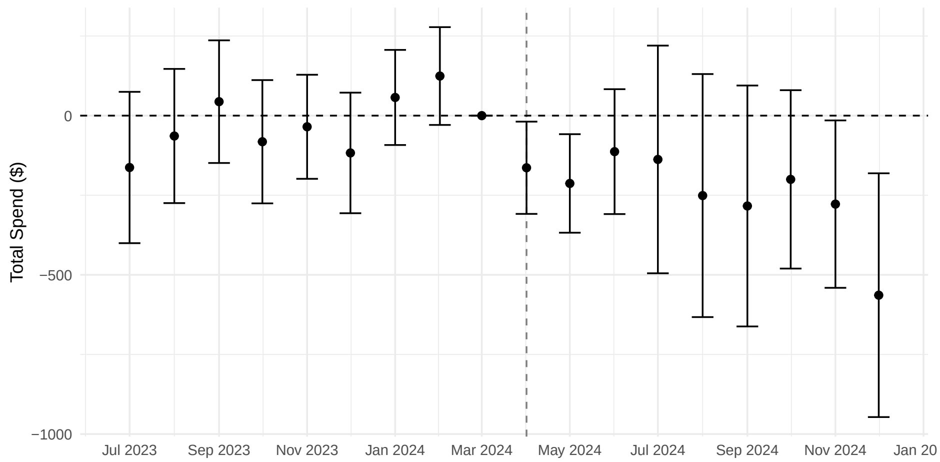

(B) Excluding Intermediate Annulus

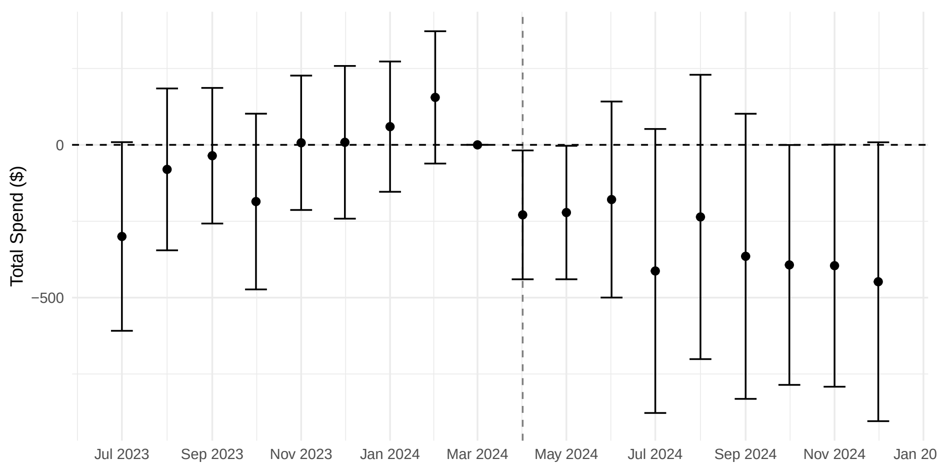

Notes. Monthly event study coefficients for total restaurant spending. Treatment is defined as restaurants within 30 minutes driving time of the Francis Scott Key Bridge. Panel (A) includes all restaurants (control group: restaurants 30–60 minutes from the bridge); Panel (B) excludes restaurants in the intermediate annulus (30–45 minutes from the bridge), using only restaurants beyond 45 minutes as controls. Vertical dashed line indicates the first month after the bridge collapse (April 2024). Pre-period coefficients cluster around zero, supporting parallel trends. Bars show 95% confidence intervals with restaurant-clustered standard errors. Data from SafeGraph Spend.

Table 1 — Effect of Bridge Collapse on Restaurant Spending

	3-Month Window		9-Month Window	
	Int. Annulus Included	Int. Annulus Excluded	Int. Annulus Included	Int. Annulus Excluded
	(1)	(2)	(3)	(4)
Bridge Collapse × Post	-223.70^{**}	-281.50^{*}	-218.69^{*}	-278.68^{*}
	(83.78)	(131.58)	(94.55)	(132.38)
Restaurant FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Observations	15,774	10,104	47,322	30,312

Notes. Two-way fixed effects DiD estimates from equation (1). Treatment defined as restaurants within 30 minutes driving time from Francis Scott Key Bridge. "Int. Annulus Incl." uses all restaurants 30–60 minutes as controls; "Int. Annulus Excl." drops restaurants in the intermediate 30–45 minute zone. 3-month window: January–June 2024; 9-month window: July 2023–December 2024. The structural estimation targets column (3). Standard errors in parentheses, clustered at the restaurant level. Data from SafeGraph Spend. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

and window durations, and as above, correspond to average declines of 4%-5% in pre-period restaurant revenue. As a further placebo check, we run the same specification using Washington, DC restaurants as treated units; these are located roughly 40 miles from the bridge and should be entirely unaffected by its collapse. The DC placebo yields no significant effect, with a null positive point estimate, confirming that the spending decline is specific to bridge proximity rather than a broader regional trend (Appendix B.10).

Lastly, while the above figures provide clear evidence that the bridge collapse reduced nearby restaurant revenue, one may be concerned that it is the first-order disruption from the disaster, rather than changes in travel time to restaurants, that drives the observed effect. In response to this, we first note that the "disaster effect" explanation is most likely to hold closest to the time of the bridge collapse, but we document very similar effects over the full 9-month post-period, and effects in these later months are unlikely to be caused by any one-time disruptions stemming the event of the collapse; if anything, point estimates are slightly smaller right after the bridge collapse, suggesting that any transient effects from the disaster itself may have attenuated rather than driven the observed estimates, for example if disaster workers provide demand for nearby restaurants. Second, we address this empirically by re-estimating the DiD after excluding all restaurants within 5 miles straight-line distance of the bridge, removing those most plausibly exposed to direct disruption, and recover a slightly larger effect ($-\$239$, $p = 0.014$, versus baseline $-\$219$). Taken together, this evidence suggests that effects are most plausibly driven by the travel time increase, both from the direct route removal and from the resulting widespread congestion (Chang et al. 2025), rather than localized bridge-site disruption. For more details on this robustness analysis, see Appendix B.11.

4.2 Consumer Demand Model

The reduced-form evidence establishes that the bridge collapse causally reduced restaurant spending, but we note that the binary treatment definition (within 30 minutes versus beyond) does not directly translate to any specific structural preference parameter β^d, as the DiD coefficient captures the average effect on treated restaurants relative to controls, but treated restaurants experienced heterogeneous travel time changes and even control restaurants may have been partially affected. We therefore perform indirect-inference estimation, following Gourieroux, Monfort, and Renault (1993), Andrews, Gentzkow, and Shapiro (2020), and Gechter and Tsivanidis (2025), to leverage our reduced-form estimate to credibly identify travel time disutility in a structural model of consumer restaurant choice.

First, we specify a parsimonious logit model in which consumer i residing in census block group (CBG) g chooses among restaurants $j \in \mathcal{J}_{gt}$ in period t to maximize utility

$$u_{ijt} = \alpha + \beta^d \cdot \log(d_{gj}) + \varepsilon_{ijt} \tag{2}$$

where α is a constant, d_{gj} is the driving time in minutes from CBG g to restaurant j, and ε_{ijt} is an i.i.d. Type I Extreme Value taste shock with an outside option normalized to zero;

see Appendix E for the full specification.[7] We define the consumer's consideration set $\mathcal{J}_{gt}$ to include all restaurants within 120 minutes driving time of CBG g.[8]

Second, to credibly recover β^d, we construct a moment condition that requires the structural model to exactly reproduce in simulation the reduced-form DiD coefficient estimated from the data, allowing us to directly leverage the causal variation from the bridge collapse to pin down β^d while accounting for heterogeneous treatment intensity and nonlinear demand responses.

Specifically, we estimate the model with two moment conditions that exactly identify the parameters (α, β^d). The first moment condition sets mean pre-period predicted transactions equal to mean observed transactions:

$$g_1(\alpha, \beta^d) = \frac{1}{J} \sum_{j=1}^{J} \left[\hat{y}_j^{\text{pre}}(\alpha, \beta^d) - y_j^{\text{pre}} \right] = 0 \tag{3}$$

which identifies α. For each candidate β^d, this uniquely determines $\alpha(\beta^d)$ via one-dimensional optimization, effectively concentrating α out. Given this $\alpha(\beta^d)$, the second moment condition matches the model-implied DiD coefficient to the empirical estimate from Table 1:

$$g_2(\beta^d) = \widehat{\text{DiD}}^{\text{model}}(\alpha(\beta^d), \beta^d) - \widehat{\text{DiD}}^{\text{data}} = 0 \tag{4}$$

where $\widehat{\text{DiD}}^{\text{model}}$ is computed by running the same binary-treatment regression on model-predicted spending, using monthly congested travel time matrices for each post-collapse

7. The model is intentionally sparse. The bridge collapse generates one dimension of exogenous variation, changes in driving time, which we use to identify β^d via indirect inference on the DiD; we leave out other dimensions to ensure that only this credibly-estimated relationship drives our later model-based demand projections. As our goal is to characterize the impact of incorporating endogenous population dynamics, we expect that this is a conservative choice, since a richer model would likely only increase the measured importance of location in determining restaurant demand. The 9-month duration of the bridge closure further limits scope for confounding, as area demographics are effectively fixed over this horizon.

8. The implied outside-option share at the estimated parameters is approximately 85–93% per decision occasion, varying with each CBG's average driving time to its consideration set, consistent with most consumer-day decision occasions not resulting in a restaurant visit; the inside-option share calibrates to total observed monthly restaurant spending.

Table 2 — Consumer Demand Model

	Estimate
Travel Time (β^d)	-0.39^{**}
	(0.149)
Constant (α)	-8.91^{***}
	(0.543)
N restaurants	2,629
N observations	47,322

Notes. β^d estimated via indirect inference, matching the model-implied DiD coefficient to the empirical estimate from Table 1. Travel time enters in logs. Bootstrap standard errors in parentheses from Bayesian bootstrap with 100 replications, resampled at the restaurant level. Data from SafeGraph Spend. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

month.[9]

4.2.1 Results

Table 2 reports the structural estimates. The bridge-identified travel time elasticity is $\hat{\beta}^d = -0.39$, corresponding to a 3.9% decline in demand for every 10% increase in travel time distance. This disutility is comparable in magnitude to the estimated retail demand elasticity over travel time in Cao et al. (2024), and implies that consumer demand is moderately concentrated in restaurants' immediate neighborhoods.

For comparison, Table C1 in the appendix reports an alternative estimate that replaces the DiD moment with a cross-sectional orthogonality condition, requiring only that spending prediction errors be uncorrelated with log travel duration (see Appendix C.1 for details). The resulting estimate is $\hat{\beta}^d = 2.08$, which has the opposite sign from the DiD-identified estimate, implying, implausibly, that longer travel time increases demand. This reversal likely

9. The structural model predicts visitors to each restaurant; we convert to spending by multiplying predicted visitors by each restaurant's observed mean spend per transaction, $\bar{m}_j = \sum_t \text{spend}_{jt} / \sum_t \text{transactions}_{jt}$, so that the model-implied DiD is in the same units as the reduced-form estimate from Table 1. We verify that this per-restaurant multiplier is approximately stable across the bridge collapse: a two-way-FE DiD (restaurant + month) of median spend per transaction on the treatment $\times$ post interaction yields an insignificant coefficient of -0.45 (SE 0.33, $p = 0.17$), about -2% of the sample mean. Post-collapse travel time changes are computed by removing the bridge from the OSRM network and augmenting with the congestion multipliers from Chang et al. (2025). See Appendix E for details.

reflects that higher-quality restaurants attract customers from a wider geographic catchment, and consumers willingly travel farther for restaurants that better match their preferences, generating a positive cross-sectional correlation between unobserved restaurant quality and average travel duration that overwhelms the true negative causal effect of travel cost.[10] The DiD-identified estimate avoids this confound by using clean within-restaurant variation in travel time induced by the bridge collapse.

5 Household Residential Choice

We next turn to household residential choice. We showed above that consumer demand, taking population geography as given, is importantly determined by preferences over travel distance; but long-term consumer demand depends also on this population geography itself, which may change endogenously from year to year. This endogenous sorting may have important implications for which areas are better able to support business locations in the long term, and may also bias observational estimates of underlying preferences over area amenities.

5.1 Identification: The Baltimore County Zoning Review

We address the identification challenge using Baltimore County's Comprehensive Zoning Map Process (CZMP), a county-wide rezoning cycle that operates on a fixed four-year schedule set by County Code, which provides sharp, exogenously-timed variation in restaurant density across years. We focus on the 2016 rezoning cycle, as this occurs after the ACS 5-year data became fully available in all fields in 2013 and before the COVID-19 period. Petitions for this cycle opened in September 2015 and then were finalized in 2016, with a substantial number of parcels newly classified into restaurant-permitting zones, while other parcels lost such classification; for more details, see Appendix B.2.1.

10. Additionally, restaurants may tend to locate nearest to their most marginal customers rather than their most loyal ones.

While this process provides sharp variation in restaurant density, we acknowledge that areas that were rezoned were almost certainly non-random, as they had to be selected for rezoning. However, we argue that the policy timing is effectively exogenous to potential confounding variation, since the four-year cycle is set by statute and does not respond to short-run market conditions. As such, we expect that any persistent confounds in terms of which areas are selected would also show up in pre-trend estimates, or in other words, that pre-trends should provide a clear test for whether or not confounded variation drives any observed results. If pre-trends are statistically null while post-period effects show significant shifts, we interpret this as a credible estimate of the effect of sharp changes in rezoning on the respective outcome. (We also present extensive robustness checks, discussed in more detail below.)

As with the Francis Scott Key bridge collapse in the preceding section, we estimate effects using a differences-in-differences model around the focal event study.

First, for the first-order effects of rezoning on restaurant entry, we estimate effects using the following equation:

$$\text{NumRest}_{gt} = \gamma_g + \delta_t + \beta^{\text{rest}} \cdot \mathbf{1}_{\text{Treat}_g \times \text{Post}_t} + \varepsilon_{gt} \tag{5}$$

where $\text{Post}_t = \mathbf{1}\{t > 2016\}$ indicates the post-rezoning period, treatment is defined as a binary indicator of a positive increase in the proportion of area zoned for restaurants, and γ_g and δ_t are CBG and time fixed effects, respectively.

Second, for the second-order effects of rezoning on population changes, we construct a CBG-level measure of nearby restaurant-zoned land by computing the share of land area within 0–1-mile, 1–2-mile, and 2–3-mile annuli around each CBG centroid that is zoned for restaurant use, on the premise that households are likely to have preferences over the number of restaurants within walking or driving distance. The focal variable is the change in the share of area zoned for restaurants between 2012 and 2016:

$$\Delta Z_g^{a\text{-}b} = \frac{\text{Area}(A_g^{a,b} \cap \mathcal{Z}_{2016})}{\text{Area}(A_g^{a,b})} - \frac{\text{Area}(A_g^{a,b} \cap \mathcal{Z}_{2012})}{\text{Area}(A_g^{a,b})} \tag{6}$$

where $A_g^{a,b}$ is the annulus between a and b miles from CBG g's centroid and $\mathcal{Z}_t$ is the union of restaurant-permitting zone polygons in year t. We then estimate the following equation:

$$\begin{aligned} \log(\text{Pop}_{gt}) = \gamma_g + \delta_t &+ \beta^{0\text{-}1} \cdot \Delta Z_g^{0\text{mi}-1\text{mi}} \times \text{Post}_t \\ &+ \beta^{1\text{-}2} \cdot \Delta Z_g^{1\text{mi}-2\text{mi}} \times \text{Post}_t \\ &+ \beta^{2\text{-}3} \cdot \Delta Z_g^{2\text{mi}-3\text{mi}} \times \text{Post}_t + \varepsilon_{gt} \end{aligned} \tag{7}$$

where $\text{Post}_t = \mathbf{1}\{t > 2016\}$ indicates the post-rezoning period.[11] For event studies, we estimate an analogous equation with indicators for each pre-period and post-period year to examine effects both prior to and after treatment.

Figure 3 presents the reduced-form event studies. For both restaurant entry and population responses, we see tightly-estimated null effects prior to the zoning change. In Panel (A), we show that zoning changes lead to significant increases in restaurant entry, with flat pre-trends and post-2016 increases peaking at roughly 0.5 additional restaurants per treated CBG in 2017 and remaining elevated through 2019.[12] Panel (B) shows the population response, which reveals a non-monotonic spatial pattern: CBGs exposed to more restaurant zoning within 1 mile experienced significant population declines starting in 2018, consistent with immediate

11. For both of the above models, we estimate over CBGs from across the Baltimore metropolitan area (3,106 CBGs observed over 2013–2019), with Baltimore County CBGs as treated units and surrounding counties as controls. However, we note that treatment intensity is continuous, and that of 526 Baltimore County CBGs, 32 experience zero change in restaurant-zoned land share across all distance rings and thus serve as additional controls.

12. We also find null effects of rezoning on non-restaurant retail establishment counts ($p = 0.35$), implying that rezoning caused restaurant entry but not non-restaurant retail entry and that non-restaurant retail entry is unlikely to have driven the observed second-order population effects. In any case, our structural household model separately includes non-restaurant retail density as an auxiliary control (Section 5), so any residual role of non-restaurant retail in household location choice is accommodated downstream.

Figure 3 — Effect of Rezoning on Restaurants and Population

Notes. Panel (A): Coefficients for restaurant count; treatment is binary indicator for zoning exposure within 1 mile. Panel (B): Coefficients for log population by distance annulus; treatment is continuous zoning exposure. Reference year is 2015 (hollow circle). Vertical dashed line indicates the 2016 CZMP vote. Pre-period coefficients are flat, supporting parallel trends. Bars show 95% CIs with CBG-clustered SEs. Data from Data Axle (restaurants) and ACS (population).

Table 3 — Effect of Restaurant Zoning on Population

	Log Population		
	0–1 mi	1–2 mi	2–3 mi
Effect of zoning	$-5.31^{\dagger}$	10.19^{*}	-7.31
	(2.78)	(4.65)	(5.91)
CBGs	3,094	3,094	3,094
CBG FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes

Notes. Two-period DiD estimates of equation (7), where the outcome is log population from ACS and treatment is the continuous change in restaurant-zoned land share within each distance annulus between 2012 and 2016 Baltimore County zoning maps. Pre-period is 2013–2016; post-period is 2017–2019. Sample includes all 3,094 nonzero-population CBGs across the 10-county study area. Standard errors in parentheses, clustered at the CBG level. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

proximity being a disamenity due to noise and traffic congestion, but CBGs with increased zoning exposure in the 1–2 mile ring experienced significant population growth starting in 2017, suggesting that restaurants at moderate distance are a positive amenity. At 2–3 miles, the rezoning produces null effects. At the average rezoning intensity in our sample, the population effect coefficient estimates imply roughly a 0.2% population decline within 1 mile and a 0.2% population increase in the 1–2 mile ring following positive restaurant rezoning, realistically modest, but also meaningful shifts that then compound over time and aggregate across neighborhoods.[13]

Table 3 reports the pooled two-period DiD estimates.[14] The 1–2 mile amenity effect is significant at the 5% level, while the 0–1 mile disamenity effect is significant at the 10% level in the pooled specification. That said, we highlight that both the 0-1mi and 1-2mi population responses are significant at the 5% level in later post-treatment years, even though this pooled estimate is moderately noisier.

13. We also note that as the ACS is based on a moving average of population, results may be best interpreted as conservative lower-bounds on true effects year-to-year.

14. Following the reduced-form event-study evidence showing that effects significantly manifest from 2017 on, we define the post-period as 2017-2019.

5.1.1 Robustness

We subject these results to several robustness tests, summarized briefly here, with full details in Appendix B.

First, we note that the non-monotonic pattern itself already narrows the scope for confounding, since most generic confounders that might correlate with rezoning, such as broader commercial development or infrastructure investment, would likely predict monotonically declining or increasing effects with distance rather than a sign reversal at 1 mile.

Nonetheless, one may be concerned that zoning changes coincide with development activity in certain corridors, and/or that time-varying, area-specific trends could confound the population response. We identify three area corridors with active commercial investment during the same period and show that excluding each one individually, or all three simultaneously, leaves point estimates essentially unchanged or if anything larger in magnitude (Appendix B.3). We also perform a more granular, data-driven version of this test, using DBSCAN spatial clustering to identify tight clusters of high-treatment-intensity CBGs and then re-estimating after dropping each cluster one at a time; this confirms that no single geographic area drives the results, with point estimates stable across all exclusions (Appendix B.4).

Additionally, one may ask whether the non-monotonic pattern is an artifact of the particular 1-mile cutoff used to define the distance annuli. We show that perturbing the annulus boundaries produces stable results across all specifications (Appendix B.5). One may also worry that the 0–1 mile disamenity reflects restaurants literally displacing residents from the same parcel or block rather than a broader neighborhood-level effect. We show that excluding own-CBG zoning from the 1-mile buffer, and separately excluding the entire inner 0.25 miles, leaves the disamenity coefficient essentially unchanged, confirming that the effect operates at the neighborhood scale (Appendix B.6). Lastly, because Baltimore County wraps around Baltimore City, one may worry that treated CBGs near the county–city boundary reflect city-specific trends rather than restaurant-specific effects. We show that excluding CBGs within 0.25 and 0.5 miles of the boundary leaves the 1–2 mile amenity unchanged and if anything sharpens the 0–1 mile disamenity (Appendix B.7).[15]

That said, while these checks probe geographic concentration, boundary sensitivity, and annulus specification, we cannot rule out all possible sources of bias, and rezoning is not as cleanly exogenous as an unexpected bridge collapse. We therefore supplement these checks with a direct downstream robustness analysis in Appendix D.3, where we re-estimate the entry model under severe attenuation and inflation of the household restaurant proximity coefficients (scaling by 0.5 and 1.5) and show that the endogenous-market improvement in total out-of-sample MSE is broadly unchanged, so our broader results would be robust even to very severe bias in these restaurant-amenity parameter estimates. We also note that our baseline household specification carries a rich set of auxiliary controls, but the entry-model improvement from incorporating endogenous population dynamics is nearly identical when we drop all of these and re-estimate with only restaurant proximity, moving cost and distance terms (Section 6.1.2), indicating that the endogenous-market advantage is not driven by the auxiliary covariates but by the DiD-identified amenity effects.

5.2 Household Location Choice Model

Motivated by the above reduced-form evidence, we specify a simple structural model of amenity-sensitive household location choice.

Each period, households decide whether to remain in their current CBG or to relocate. A household in CBG g obtains utility $U_{ggt} = V_{gt} + \varepsilon_{ggt}$ from staying, and evaluates destination $g' \neq g$ according to $U_{gg't} = \beta^{\text{move}} + V_{g't} + \beta^{\text{dist}} \cdot d_{gg'} + \varepsilon_{gg't}$, where $d_{gg'}$ is the log driving time between CBG centroids. The key parameters are β^{move} (baseline moving cost), β^{dist} (distance-related moving cost), and the flow utility from location g':

$$V_{gt} = \beta_1^{\text{rest}} \cdot R_{gt}^{1\text{mi}} + \beta_2^{\text{rest}} \cdot R_{gt}^{2\text{mi}} + \mathbf{X}'_{gt}\boldsymbol{\gamma} + \xi_g \tag{8}$$

15. We also verify that the population responses are not driven by rent-based displacement. Re-estimating with log median rent as the outcome, we find that rents, if anything, move in the same direction as the population effect, ruling out a story in which rent increases push residents out (or rent decreases pull residents in). See Appendix B.8.

where $R_{gt}^{1\text{mi}}$ and $R_{gt}^{2\text{mi}}$ are counts of restaurants within 1 mile and within 2 miles of the CBG centroid, inspired by our empirical evidence showing that restaurant density changes in these annuli produce significant effects (and that density changes in the 2–3 mile annulus produces null effects, so we leave that term out); $\mathbf{X}_{gt}$ is a vector of neighborhood covariates entering with coefficient $\boldsymbol{\gamma}$; and ξ_g captures unobserved location quality.[16] We estimate two specifications: a "baseline" in which $\mathbf{X}_{gt}$ includes non-restaurant business density in each annulus, log local income and log population density,[17] distance to the Baltimore CBD (with a DC-corridor interaction), and a same-county fixed effect; and a parsimonious "no-controls" specification in which $\boldsymbol{\gamma}$ is set to zero. The baseline residualizes local restaurant density against these correlated neighborhood characteristics; the no-controls specification serves as a downstream robustness check confirming that the restaurant-proximity estimates are not driven by the specific choice of controls (see Appendix D.2).[18]

5.2.1 Estimation

We then fit this model of household location choice via GMM, using indirect inference to directly leverage our quasi-experimental variation in identifying our structural estimates of restaurant amenity preferences (Gourieroux, Monfort, and Renault 1993; Gechter and Tsivanidis 2025; Andrews, Gentzkow, and Shapiro 2020).

Specifically, given candidate parameters $\theta = (\beta^{\text{move}}, \beta_1^{\text{rest}}, \beta_2^{\text{rest}}, \beta^{\text{dist}}, \boldsymbol{\gamma})$, we simulate popu-

16. We do not separately recover the unobserved component ξ_g. Because the reduced-form DiD regressions that discipline the restaurant coefficients include CBG fixed effects, persistent location quality ξ_g is differenced out at identification and cannot confound the restaurant proximity estimates; its absolute level is absorbed into the 2013 baseline at which we initialize the simulation, so the restaurant proximity coefficients are accordingly the marginal valuation of restaurant access, conditional on these unobserved factors.

17. Log income and log population density each enter with two coefficients: an origin ("push") coefficient on the value at the household's current CBG and a destination ("pull") coefficient on the value at each candidate destination. This asymmetry is needed to fit the county-flow moments g_5^{od}, which show high-income, high-density CBGs (e.g., urban cores) simultaneously exhibiting both high inbound and high outbound mover rates: a specification with only destination attractiveness would predict monotonic accumulation toward such areas rather than the observed churn. Full estimates are in Appendix D.2.

18. Zoning reform does not significantly move these auxiliary covariates. Appendix B.9 shows that replacing log population with log median income or education share (BA+) as the outcome in equation (7) yields null effects across all distance rings, so their omission (as in the no-controls specification) is unlikely to confound the restaurant proximity estimates.

lation dynamics across all 3,106 CBGs over 2013–2019 by iterating the population evolution equation from the base of 2013 population, then compute moments from the simulated data. The two identifying moments for the restaurant-proximity coefficients are the indirect-inference conditions built from the reduced-form DiD:

$$g_1(\theta) = \hat{\beta}^{\text{sim,0-1}}(\theta) - \hat{\beta}^{\text{data,0-1}} = 0 \tag{9}$$

$$g_2(\theta) = \hat{\beta}^{\text{sim,1-2}}(\theta) - \hat{\beta}^{\text{data,1-2}} = 0 \tag{10}$$

which match the coefficients of the reduced-form DiD regression (7) on model-simulated population movement to the empirical DiD estimates from Table 3. These moments carry the quasi-experimental variation into the structural model, and are the primary source of identification for β_1^{rest} and β_2^{rest}.[19]

The remaining structural parameters (the moving cost β^{move}, the distance decay β^{dist}, and the vector of neighborhood covariate coefficients $\boldsymbol{\gamma}$) are disciplined by two aggregate Census ACS moment conditions on mobility patterns and a vector of county-to-county flow moments:

$$g_3(\theta) = \text{SameCountyRate}^{\text{sim}}(\theta) - 0.761 = 0 \tag{11}$$

$$g_4(\theta) = \text{MobilityRate}^{\text{sim}}(\theta) - 0.0968 = 0 \tag{12}$$

$$g_5^{od}(\theta) = s_{od}^{\text{sim}}(\theta) - s_{od}^{\text{data}} = 0 \quad \text{for each origin-destination pair } (o, d) \tag{13}$$

g_3 matches the proportion of moves within the same county (76.1%). g_4 matches the overall annual mobility rate (9.68%). g_5 is a vector of 84 county-to-county flow moments, one for each observed origin-destination pair (o, d), where s_{od} is the fraction of movers leaving county o

19. Strictly, the DiD moment g_1 (analogously g_2) equates the simulated to the empirical population response to zoning, which in the structural model reflects both restaurant preferences β^{rest} acting on zoning-induced restaurant entry and non-restaurant retail preferences γ^{nonrest} acting on any zoning-induced non-restaurant retail entry. In practice, since we estimate null effects of zoning on non-restaurant retail entry, this DiD moment is expected to load primarily on β^{rest}; the additional moments below (mobility, same-county, and county-flow) discipline γ^{nonrest} against the observational cross-sectional variation, separating it from β^{rest}.

who move to county d, derived from Census ACS county-to-county flows; these flow moments discipline $\boldsymbol{\gamma}$ by requiring the sorting model to reproduce observed cross-county sorting on income, density, and other neighborhood characteristics. For further details, see Appendix E.

Letting $\mathbf{g}(\theta) = (g_1, g_2, g_3, g_4, g_5)'$, we minimize the weighted objective:

$$Q(\theta) = \mathbf{g}(\theta)'\mathbf{W}\,\mathbf{g}(\theta) \tag{14}$$

where $\mathbf{W}$ is a diagonal weighting matrix; see Appendix E for details.

5.2.2 Results

Table 4 reports the focal parameter estimates (full estimates in Appendix D.2). The restaurant amenity parameters align closely with the observed non-monotonic pattern in the reduced-form data, as restaurants within 1 mile are a disamenity ($\hat{\beta}_1^{\text{rest}} = -0.038$) while restaurants at 2-mile proximity are an amenity ($\hat{\beta}_2^{\text{rest}} = 0.023$). The 2-mile amenity coefficient is significant at the 0.1% level and the 1-mile disamenity at the 1% level; the joint Wald test rejects the null that both restaurant proximity parameters are zero at $p < 0.001$. The maximally-sparse household specification recovers qualitatively similar restaurant-proximity coefficients; see Appendix D.2. Since $d_{gg'}$ enters in log driving time, we can express the amenity magnitude in driving-time-equivalent units (via $\hat{\beta}^{\text{dist}} = -4.25$ per log-time unit): ten additional restaurants in the 1–2 mile ring raise CBG attractiveness by the same amount as roughly a 5% reduction in moving distance.

As with the consumer demand model, we compare the quasi-experimental and observational household model estimates. Without the DiD moments from the zoning shock, the point estimates change substantially; the moving cost grows in magnitude from −6.71 to −10.10, and the 1 mile disamenity grows by roughly a factor of four in normalized scale (from 0.6% to 2.5% of $|\beta^{\text{move}}|$), while the 2 mile amenity moves in the opposite direction, shrinking to about a third of its DiD-identified normalized magnitude. This pattern is consistent with observational moments capturing equilibrium correlations between population and restaurant

Table 4 — Household Location Choice Model Estimates

Parameter	Estimate
Restaurant proximity, 1 mi (β_1^{rest})	-0.038^{**}
	(0.011)
Restaurant proximity, 2 mi (β_2^{rest})	0.023^{***}
	(0.002)
Moving distance decay (β^{dist}, per 10 mi)	-4.25^{***}
	(0.003)
Other auxiliary controls	Yes
Joint Wald (restaurant params = 0)	$p < 0.001$
N CBGs	3,106

Notes. Estimates from GMM via indirect inference, matching simulated moments to empirical targets including DiD moments from the zoning change. Restaurant proximity coefficients measure the effect of restaurant count within each distance ring on CBG utility. Other auxiliary controls (included in estimation but suppressed from display) are the moving cost, distance-to-CBD, same-county fixed effects, non-restaurant business density in the same annuli, and log local income and population density push/pull terms; full parameter estimates are reported in Appendix D.2. Bootstrap standard errors in parentheses from Bayesian bootstrap with 100 replications, resampled at the CBG level; each replication recomputes the DiD targets from reweighted data, so the standard errors account for first-stage sampling uncertainty. See Appendix E.4 for details. Data from ACS. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

density rather than the causal effect of restaurant density on residential location choice. For full details, see Appendix C.2.

6 Restaurant Entry

The above analyses furnish us with two well-identified structural models that capture 1) how contemporaneous consumer demand varies with restaurant proximity and 2) how population adjusts over a longer time horizon in response to restaurant amenities. With these in hand, we now examine implications for restaurant location choice.

6.1 Dynamic Entry Model

We model entry and exit decisions across all 3,106 census block group (CBG) centroids in the Baltimore metropolitan area, matching the spatial resolution of the upstream consumer demand and household location models.[20] At each location j and period t, the market contains n_{jt} active incumbent restaurants and a pool of $K = 100$ potential entrants, indexed by i. Each firm i chooses an action a_{it} at the end of period t: an incumbent's choice is $a_{it} \in \{\text{cont}, \text{exit}\}$; a potential entrant's choice is $a_{it} \in \{\text{enter}, \text{not-enter}\}$. The location-period state is $x_{jt} = (n_{jt}, M_{jt})$, where M_{jt} aggregates expected visits from each origin CBG using the demand model estimates from Section 4. The model nests both exogenous-market and endogenous-market specifications, differing only in how M_{jt+1} is forecast.

An active restaurant earns per-period flow profits

$$\pi_{jt} = \mu \cdot \frac{M_{jt}}{n_{jt}} \tag{15}$$

where μ scales per-restaurant demand revenue under equal sharing among symmetric co-located competitors.[21] Each firm i draws a private idiosyncratic shock $\varepsilon_{it} \sim N(0, 1)$ at the time of its action, observed only by that firm and independent across firms and time; the aggregate state (n_{jt}, M_{jt}) is common knowledge. Note that, with the standard unit-variance normalization, cost and value parameters are here denominated in units of the shock ε_{it}.

We follow the moment-based Markov equilibrium (MME) framework of Ifrach and Weintraub (2017), which generalizes the oblivious equilibrium of Weintraub, Benkard, and Van Roy (2008); firms condition on the aggregate state (n_{jt}, M_{jt}) rather than on rivals' individual states, and the aggregate transition $F(x_{j,t+1} \mid a_{it}, x_{jt})$ treats firms in n_{jt} symmetrically. This aggregation is exact in the large-market limit and a standard approximation otherwise. An incumbent's continue/exit choice contributes to $n_{j,t+1}$ as one of n_{jt} symmetric draws, while an entrant's 1 contribution to $n_{j,t+1}$ is treated separately. See Appendix E for full details.

20. CBGs are a standard Census geographic unit; this spatial resolution naturally matches the CBG-level consumer origins used in the demand model and the CBG-level population used in the household model. We restrict candidate locations to CBGs with at least one restaurant present at any point during our sample period to filter out areas not zoned for commercial activity.

21. We note that the substantive spatial-competition channel operates upstream in the demand model. When a restaurant enters near j, P_{gj} falls through the logit denominator and M_j shrinks accordingly. The division by n_{jt} in equation (15) just allocates M_j among restaurants co-located within the same CBG, a small number at this spatial resolution.

Define the choice-specific continuation value as the expected next-period integrated incumbent value conditional on staying,

$$v^{\text{cont}}(x_{jt}) = \mathbb{E}\Big[V(x_{j,t+1}) \,\Big|\, a_{it} = \text{cont},\, x_{jt}\Big]\,, \tag{16}$$

where the conditional expectation is over the transition F, with n evolving under estimated conditional choice probabilities (CCPs) and M evolving under either the structural household model or exogenous projections. Define analogously $V^{\text{ent}}(x_{jt}) = \mathbb{E}[V(x_{j,t+1}) \mid a_{it} = \text{enter}, x_{jt}]$ for the entrant. Under a profit-then-decide timing convention, the firm receives flow profit at the start of period t and chooses its action at the end of the period, so the action-specific values satisfy

$$v^{\text{stay}}(x_{jt}) = \pi_{jt} + \delta\, v^{\text{cont}}(x_{jt}), \qquad v^{\text{exit}}(x_{jt}) = \pi_{jt} + \delta\, SV, \tag{17}$$

where $\delta = 0.95$ is the discount factor and SV is the scrap value received upon exit. We define $V(x_{jt})$ as the integrated continuation value under the equilibrium CCP policy. Following Aguirregabiria and Mira (2007), the integrated value satisfies

$$V(x_{jt}) = \sum_a P(a \mid x_{jt}) \cdot \Big[\pi(a, x_{jt}) + e^P(a, x_{jt}) + \delta\, \mathbb{E}[V(x_{j,t+1}) \mid a, x_{jt}]\Big]\,, \tag{18}$$

where $e^P(a, x)$ is the conditional expectation of the private shock $\varepsilon_{it}(a)$ given that action a is chosen.

Entry and exit decisions then follow probit threshold rules,

$$p^{\text{exit}}(x_{jt}) = 1 - \Phi\left(\frac{\delta\,(v^{\text{cont}}(x_{jt}) - SV)}{\sigma}\right), \qquad p^{\text{enter}}(x_{jt}) = \Phi\left(\frac{\delta\, V^{\text{ent}}(x_{jt}) - EC}{\sigma}\right), \tag{19}$$

where EC is the entry cost (paid at the time of the entry decision) and Φ is the standard normal CDF.[22]

22. Under the profit-then-decide timing convention above, π_{jt} cancels from the incumbent's choice differential, $v^{\text{stay}} - v^{\text{exit}} = \delta\,(v^{\text{cont}}(x_{jt}) - SV)$, and so it does not appear in the exit threshold. Because shocks are firm-specific, the model accommodates simultaneous entry and exit in the same location-year.

Under the endogenous-market specification, restaurants use the household model from Section 5 to forecast how entry affects future population:

$$\mathbb{E}[M_{j,t+1}] = f(M_{jt}; \hat{\beta}_1^{\text{rest}}, \hat{\beta}_2^{\text{rest}}, \mathbf{Rest}_t^{1\text{mi}}, \mathbf{Rest}_t^{2\text{mi}}) \tag{20}$$

where $f(\cdot)$ iterates the household location choice model forward one period, incorporating how cumulative restaurant density within 1 mile and within 2 miles of each CBG affects CBG utility and thus population flows. We note that the per-firm state remains (n_{jt}, M_{jt}); the metropolitan-area cumulative-count vectors $\mathbf{Rest}_t^{1\text{mi}}$ and $\mathbf{Rest}_t^{2\text{mi}}$ enter only through the forward simulation that updates $M_{j,t+1}$. Under the exogenous-dynamics specification, expectations of market size growth are instead modeled using official population projections from the Maryland Department of Planning.

6.1.1 Estimation

Let $\theta = (\mu, EC, SV)$ denote the structural parameters to be estimated. Estimation uses the nested pseudo-likelihood (NPL) algorithm of Aguirregabiria and Mira (2007), adapted to the continuous market-size state via forward simulation and polynomial sieve approximation (Bajari, Benkard, and Levin 2007). We first estimate flexible conditional choice probabilities (CCPs) for entry and exit from the data as a function of the state $(n, \log M)$, and compute the integrated incumbent value V via Monte Carlo forward simulation under these CCPs. Next, we approximate $v^{\text{cont}}(x_{jt})$ and $V^{\text{ent}}(x_{jt})$ by polynomial-sieve evaluation of V at the relevant post-action states (the post-stay state for v^{cont}, and the post-entry state for V^{ent}). We then maximize the pseudo log-likelihood over θ, matching the action-specific probit probabilities to observed decisions, and update the CCPs using the structural model's implied choice probabilities, iterating this procedure until the CCPs are consistent with the structural model's implied probabilities at the estimated parameters. See Appendix E for complete computational details.

Our two-stage approach—recovering demand and household preferences first from the quasi-experimental variation, then estimating the entry game conditional on those upstream estimates—further helps to establish this entry model's identification. Because the demand and household models are identified from credible reduced-form variation (the bridge collapse and zoning change) rather than from functional form assumptions, the market size and population forecasts entering the entry model reflect causal relationships, and entry costs and profit parameters are then identified from variation in entry and exit rates conditional on these inputs. Specifically, the profit scale μ is identified from the overall level of entry and exit rates given market size; higher μ implies larger profit flows, encouraging more entry and less exit. The entry cost EC is identified from the entry threshold, since locations with higher $V^{\text{ent}}(x_{jt})$ should attract more entrants, and the entry cost governs how high $V^{\text{ent}}(x_{jt})$ must be to induce entry. The scrap value SV is identified from the exit threshold; a large (positive) scrap value would cause incumbents to exit even moderately profitable locations to collect the recoverable payout.

As the goal of our exercise is model comparison, we split our estimation sample into a training sample, defined as 2013–2017, and a hold-out test sample, defined as 2018–2019. This allows us to compare model fit on out-of-sample prediction error, with the test sample restricted to the final years to avoid look-ahead bias (Tashman 2000).

6.1.2 Results

Results are presented in Table 5. The scrap value SV_0 is positive in both specifications, indicating that exit yields a small residual payoff on average. The entry cost EC_0 is larger in the exogenous-market specification (5.02) than in the endogenous-market specification (4.27), consistent with the endogenous-market model attributing more of the cross-sectional variation in exit and entry rates to anticipated market evolution rather than to fixed costs.

On the 2018–2019 hold-out data, the endogenous-market specification reduces out-of-sample MSE by 5.37% relative to the exogenous-dynamics baseline under our baseline household utility (1.898 vs. 2.006; Diebold–Mariano $p < 0.001$), and by 7.50% under a

Table 5 — Restaurant Entry Model Estimates

	Exogenous-Market (Department of Planning projections)	Endogenous-Market (structural HH)
Parameter Estimates		
μ (profit scale)	0.222	0.086
	[−0.218, 1.025]	[0.001, 0.155]
EC_0 (entry cost)	5.017	4.273
	[2.65, 7.63]	[3.80, 4.59]
SV_0 (scrap value)	1.627	0.833
	[−0.90, 4.39]	[0.28, 1.19]
Out-of-Sample Fit (2018–2019)		
Total MSE	2.006	1.898

Notes. Nested pseudo-likelihood estimates. The exogenous-market specification uses Maryland Department of Planning population projections to forecast market size evolution; the endogenous-market specification incorporates anticipated household sorting via the structural household model. 95% bootstrap percentile confidence intervals in brackets on structural parameters (μ, EC_0, SV_0) are from a paired Bayesian bootstrap resampled at the CBG level, propagating upstream demand and household parameter uncertainty through re-estimation of the entry model in each draw; percentile CIs rather than standard errors because the nonlinear propagation produces potentially asymmetric sampling distributions (see Appendix E.4). The exogenous-market specification's marginal intervals on μ and SV_0 are wide and cross zero, reflecting the well-known weak identification of fixed-cost-type parameters in dynamic entry games with a single binary action and a fixed shock-variance restriction; the endogenous-market specification's marginal intervals are noticeably tighter, with μ excluding zero at the 5% level and SV_0 solidly positive. Joint identification of the parameter triple is preserved in both specifications (Aguirregabiria and Suzuki 2014). Out-of-sample fit statistics are reported at the point-estimate upstream; significance of the endogenous-market MSE improvement is tested via a Diebold–Mariano test on per-CBG combined-MSE residuals, yielding $p < 0.001$. Estimation uses 3,106 CBGs. Data from Data Axle.

maximally-sparse household specification that drops all auxiliary controls (1.856 vs. 2.006; $p < 0.001$); see Appendix D.2.[23] The qualitative contribution of the endogenous-market machinery is larger still: the two specifications diverge substantially in *where* they attribute value and *which* interventions they recommend—for both a policymaker choosing where to invest transit dollars and a firm choosing where to locate—patterns we document in Sections 6.2 and 6.3.

One may be concerned that the endogenous-market advantage proxies for the endogenous-

23. This wedge is a directional lower bound on the importance of endogenous urban-sorting dynamics: our household model estimates a near-null non-restaurant retail amenity effect and our demand model covers only restaurant visits, so any additional endogeneity operating through non-restaurant retail sorting would compound rather than offset the estimated advantage.

market household model accounting for area amenities more generally, rather than the restaurant→population feedback channel specifically. We test this by re-estimating the entry model with the household restaurant proximity coefficients set to zero ($\hat{\beta}_1^{\text{rest}} = \hat{\beta}_2^{\text{rest}} = 0$), holding the remaining household coefficients at their headline estimates. This mechanism-null specification collapses the wedge from 5.37% to −0.16% (total MSE 2.009, essentially indistinguishable from the exogenous baseline's 2.006): the restaurant→population feedback channel, not household sorting more generally, drives the endogenous-market advantage. See Appendix D.1 for details.

Re-estimating under a maximally-sparse household specification (moving cost, restaurant proximity, and distance terms only) yields an even larger, highly significant endogenous-market MSE improvement of 7.50% (Diebold–Mariano $p < 0.001$; see Appendix D.2): our baseline choice to include the auxiliary controls is therefore a conservative one, and the sparse specification, if anything, sharpens rather than dilutes the endogenous-market advantage.

6.2 Comparative Patterns

Having established that the endogenous-market model outperforms on out-of-sample fit, we now examine where and how the two specifications diverge.

6.2.1 Geographic Divergence

Figure 4 maps rank divergence between the two specifications, highlighting where they disagree on relative location rankings. Areas where the endogenous-market model ranks locations relatively higher (red) are concentrated in Montgomery County (including Bethesda and Silver Spring), an amenity-dense Maryland county at the DC-adjacent southern edge of our study area, where the endogenous-market specification anticipates that the existing amenity profile will continue to attract additional households; a smaller secondary red concentration appears in the central portion of the study area. Areas where the exogenous-market model ranks locations relatively higher (blue) are concentrated in Prince George's County and in the

Figure 4 — Rank Divergence Between Endogenous- and Exogenous-Market Valuations

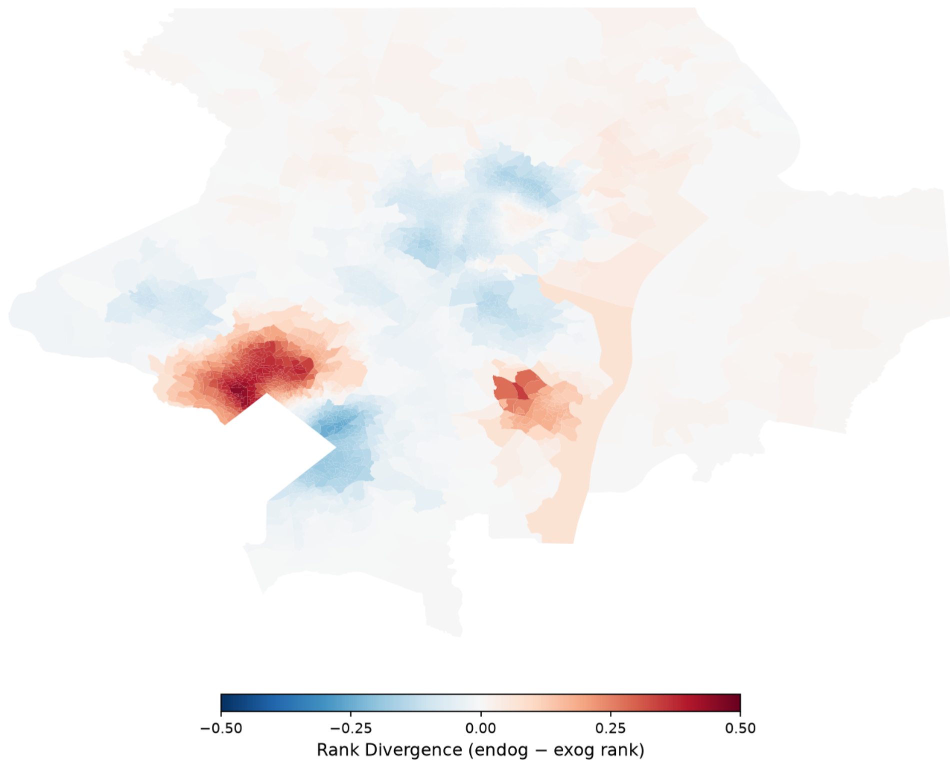

Notes. Rank divergence at the CBG level, defined as the endogenous-market percentile rank minus the exogenous-market percentile rank of each location's valuation. Red indicates locations where the endogenous-market model ranks them relatively higher; blue indicates locations where the exogenous-market model ranks them relatively higher. Valuations are spatially smoothed via a Gaussian kernel of bandwidth 2 km before percentile ranks are computed, for visual clarity. Shown across 3,106 CBGs in the 10-county Baltimore metropolitan area.

inner Baltimore County / Baltimore City area, where a static reading of current demographics would forecast stronger demand than the endogenous specification, which anticipates that current density either saturates residential preferences or generates offsetting sorting responses. Outer counties (Carroll, Harford, Howard) exhibit near-zero divergence between the two specifications. We note that this divergence reflects differences in predicted trajectories rather than in unobserved location quality; the endogenous-market model ranks these areas differently because it anticipates that restaurant density will itself feed back into household sorting over time, not because it assigns higher intrinsic value to these neighborhoods.

6.3 Counterfactual Road Analysis

The comparison above establishes that the endogenous-market model outperforms on out-of-sample prediction. We now show that the two specifications also produce substantively different counterfactual predictions for transit interventions—predictions that matter both for a policymaker choosing which corridor to prioritize and for a firm projecting how a candidate location's profitability would be affected by a transit change.

6.3.1 Corridor Comparison

We examine a stylized transit improvement, modeled as a representative 10% reduction in travel time along a hypothetical road corridor, applied to three candidate corridors selected to span a density gradient (Figure 5): (i) an "Urban Link" continuing south from Towson into north Baltimore City, with moderate existing restaurant density and substantial residential headroom, i.e. areas where additional restaurants would still fall in the 1–2 mile amenity range rather than the 0–1 mile disamenity range; (ii) an "Downtown Connector" spanning denser inner Baltimore neighborhoods, with many households but less headroom; and (iii) "Suburban Link", a corridor in the northwestern suburbs of Baltimore County that is relatively underserved in restaurants per capita compared with the two urban corridors, but more limited density in nearby areas. For each corridor we identify all origin CBGs and destination restaurant locations whose centroids fall within the corridor buffer and reduce travel times by 10% for all origin–destination pairs in which both endpoints lie within this buffer.[24]

Estimated effects on total metro-area restaurant entry are presented in Figure 6. The two specifications produce substantially reshuffled corridor rankings. The exogenous-market specification identifies "Suburban Link" as the strongest intervention (0.255% entry), followed by the "Urban Link" (0.189%), with the "Downtown Connector" the weakest (0.040%). The

24. Appendix D.4 provides corridor definitions, the full simulation procedure, and a firm-surplus decomposition for each corridor.

Figure 5 — Counterfactual Road Corridors

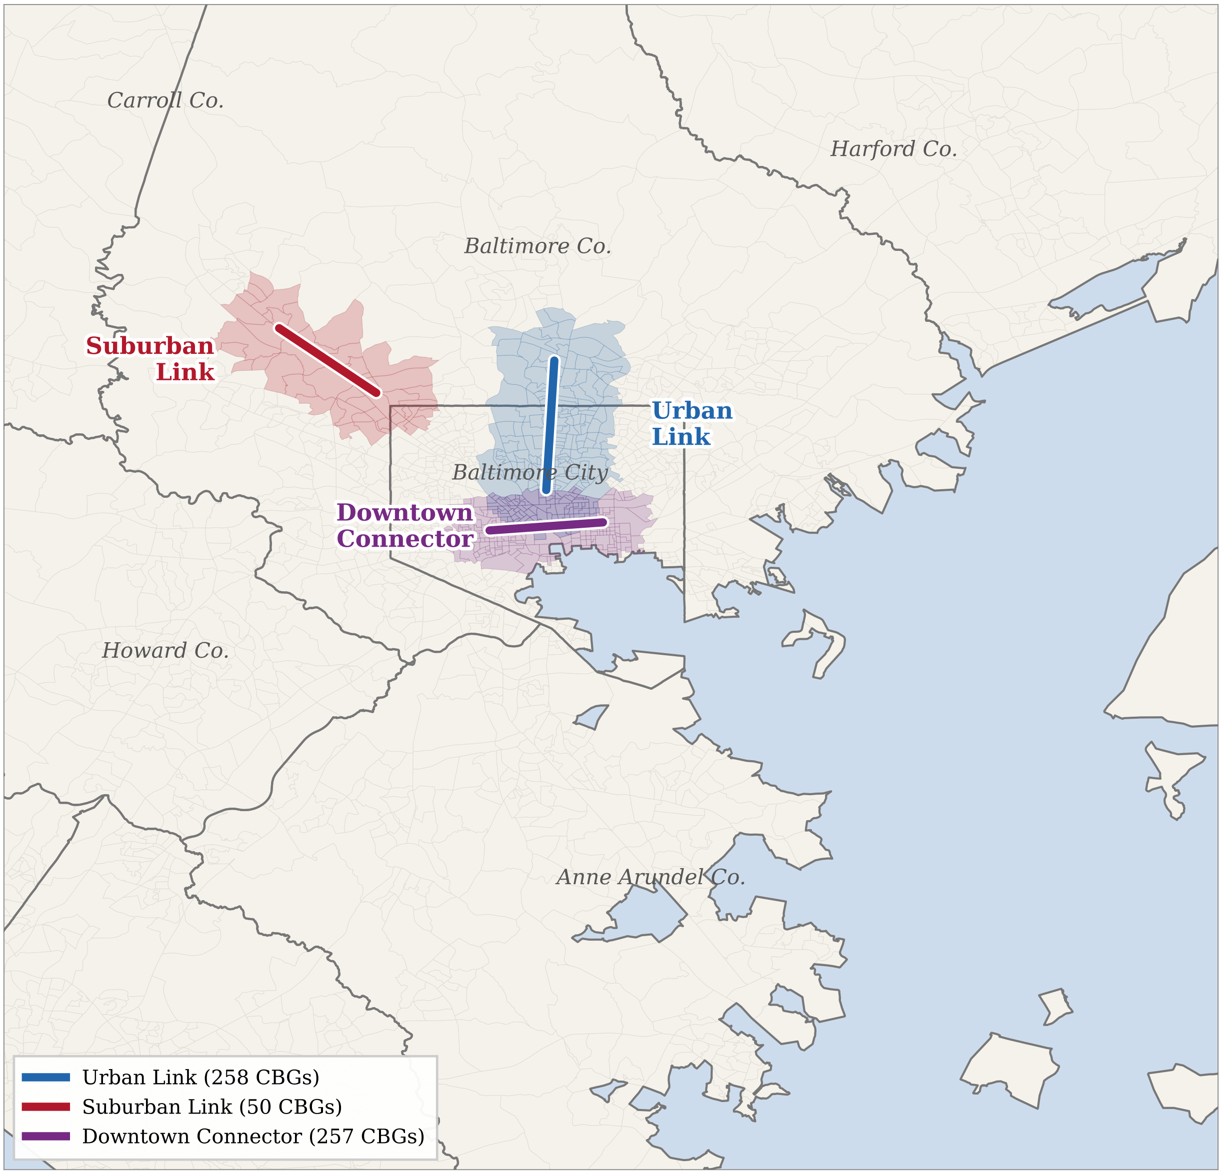

Notes. Three hypothetical road corridors spanning the Baltimore suburban ring. Shaded regions indicate CBGs whose centroids fall within the corridor buffer (1.5–2.5 miles depending on corridor scale). The travel-time reduction applies to all origin–destination pairs in which both the origin CBG and the destination restaurant location lie within the buffer. Figure 6 reports entry changes by corridor and model specification.

Figure 6 — Counterfactual Entry Response Across Road Corridors

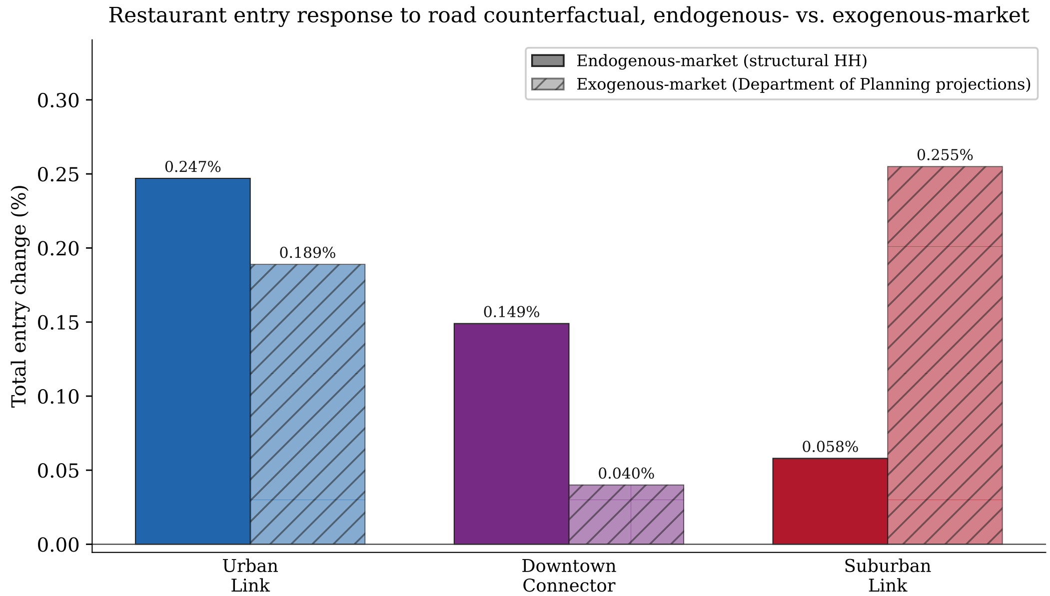

Notes. Percentage change in total metro-area restaurant entry under a counterfactual 10% travel-time reduction along each corridor, paired across the endogenous-market specification (solid bars) and the exogenous-market specification (hatched bars). Bar colors match the corridor colors in Figure 5. Per-corridor welfare decomposition is reported in Appendix D.4.

endogenous-market specification reverses the ranking, with "Suburban Link" falling from first to last: the "Urban Link" is now strongest (0.247%), the "Downtown Connector" is second (0.149%), and "Suburban Link" collapses to last (0.058%).

Under the exogenous-market specification, transit improvements raise entry by shrinking travel distance between a fixed population and existing restaurants: the model treats population as given and reads off the mechanical visit-uplift, which is positive at all three corridors. The endogenous-market specification adds a second channel, since transit-induced restaurant entry shifts area amenity density and households sort accordingly along the non-monotonic residential preference identified in Section 5.2.2, in which restaurants are an amenity at 1–2 miles but a disamenity at 0–1 miles. Whether this second-round feedback amplifies or dampens the first-round entry response depends on how much residential headroom the corridor has: the sorting channel amplifies where existing restaurant density is moderate and dampens where existing density is high, as additional entrants approach the 0–1 mile disamenity range for a larger share of already-served neighborhoods. The strength of the sorting response also scales with the local density of households within moving reach, so the amenity effect operates most strongly in corridors that combine amenity-range headroom with a large local catchment. The three corridors sit at three points along this density gradient. The "Urban Link", with moderate existing density and room for additional restaurants to land within the amenity range for most nearby CBGs, sits in the pure amplifying regime; households sort in, the market expands, and the endogenous specification exceeds the exogenous. The "Downtown Connector" sits at an intermediate point: existing density is high enough that a share of new entrants approaches the 0–1 mile disamenity range at the densest points, partially offsetting the local sorting response, so the endogenous specification still exceeds the exogenous but the absolute entry response is smaller than at the "Urban Link". "Suburban Link" currently has a low restaurant-to-population ratio, and the exogenous specification reads this as strong entry potential under the mechanical demand model: existing households are underserved, so adding restaurants captures unmet demand. The endogenous specification agrees the current population would benefit from more restaurants but sees that the compound entry response, which depends on the further sorting-in of additional households, is limited by how few nearby households there are to draw in, so aggregate entry response is bounded. A planner using the exogenous specification would rank "Suburban Link" as the stronger intervention; the endogenous specification reverses that ranking, identifying the "Urban Link" as the higher-value intervention with the "Downtown Connector" above "Suburban Link".

The three corridors trace an inverse-U in the value of endogenous dynamics. Sorting operates as a multiplier on the exogenous-market visit-uplift, and it requires two inputs to work: a nearby household base to draw in and marginal-response room on the amenity curve. The "Suburban Link" sits on the left tail, with too sparse a local catchment for sorting to multiply on; its endogenous response is modest, while the exogenous specification overstates entry by 0.197 pp. The "Downtown Connector" and "Urban Link" sit near the peak, where both inputs are present and the endogenous–exogenous gap turns positive: the "Downtown Connector", denser and with a smaller exogenous baseline, carries the largest relative wedge

(0.109 pp), while the “Urban Link”, with more amenity-range headroom, carries a smaller wedge (0.058 pp) but the largest absolute endogenous entry response. Very-mature corridors would sit on the right tail, where the marginal restaurant falls within the 0–1 mile disamenity range for enough households that the disamenity outweighs the 1–2 mile amenity benefit, though our three-corridor sample does not directly probe this range.

The same corridor-level entry numbers have a firm-level reading. Under free entry, the projected long-run entry response along a corridor indicates the long-run profitability of candidate restaurant locations within it, so the endogenous–exogenous gap at each corridor indicates the amenity-feedback correction to a naive projection of that profitability: the “Urban Link” carries a 0.058 pp positive correction (exogenous 0.189%, endogenous 0.247%), the “Downtown Connector” a 0.109 pp positive correction (0.040%, 0.149%), and the “Suburban Link” a 0.197 pp overstatement (0.255%, 0.058%). The “Downtown Connector” is where the naive projection understates profitability most severely in relative terms, while the “Urban Link” has the largest absolute endogenous projection. The direction of the correction varies across candidate corridors: positive at urban corridors where local sorting would compound over time, and negative at sparse suburban corridors where Department of Planning projections outrun the local catchment. A firm using only exogenous demographic inputs will therefore systematically mis-rank candidate corridors along this dimension.

7 Conclusion

The answer to where restaurants should locate today depends on both the present and the future. Traditional methods estimate location value based on exogenous demographics, but if households move toward neighborhoods with appealing amenities, then an influx of new restaurants can also reshape who lives nearby tomorrow. This feedback loop makes estimation exceedingly difficult, as restaurants locate where demand is high and households sort toward desirable neighborhoods, confounding relationships in observational data, but also presents an opportunity to improve on standard practices. In this paper, we develop an integrated approach, grounded in well-identified reduced-form estimation, to quantify and test this full feedback loop, developing novel implications for policymakers and restaurant managers alike.

First, we obtain credible estimates of the effect of travel time on consumer demand for restaurants using the March 2024 Francis Scott Key Bridge collapse as an exogenous shock to the surrounding transit network, establishing that consumer demand for a given restaurant declines by 3.9% for every 10% increase in travel time. Second, we obtain credible estimates of household preferences for living near restaurants using a large-scale rezoning in Baltimore County that created sharp variation in restaurant density, identifying a strong but non-monotonic residential preference for restaurants nearby, where higher immediate proximity to restaurants is a disamenity but higher intermediate proximity to restaurants is a positive amenity. We show that these estimates are robust to a wide battery of robustness checks and demonstrate that observational estimation of both consumer demand and household sorting preferences, without the quasi-experimental variation from the bridge collapse and zoning change respectively, leads to severely biased coefficients.

With these credible preference estimates in hand, we then integrate our estimated models of consumer demand and household location choice into a dynamic entry model estimated via nested pseudo-likelihood, and compare an endogenous-market specification (which anticipates entry-induced household sorting) against a benchmark exogenous-market specification using official population projections. Overall out-of-sample MSE improves by 5–8% under the endogenous-market specification, and the two specifications identify substantively different locations as high-value and rank candidate road corridors in reversed orders—with parallel implications both for policymakers choosing where to invest transit dollars and for firms projecting the long-run profitability of a candidate location in changing urban environments. We show that the advantage is driven specifically by the restaurant→population feedback channel, since setting the household restaurant proximity coefficients to zero eliminates the out-of-sample improvement. Geographically, the endogenous-market specification revalues dense urban corridors where a large local household base can compound the sorting response, while the exogenous-market specification favors suburban corridors where Department of Planning projections are optimistic but the local catchment is sparse.

Where the two specifications most consequentially diverge is in counterfactual predictions. Across a stylized transit-improvement counterfactual applied to three Baltimore-area corridors spanning a density gradient, the endogenous- and exogenous-market specifications reverse the top and bottom of the ranking: the exogenous specification favors a sparse-catchment suburban corridor whose Department of Planning projections predict strong future growth, while the endogenous specification favors a moderate-density corridor with a large local household base that can compound the sorting response, since aggregate entry in the sparse corridor is limited by how little of the local population can respond to the transit-induced amenity change. Those rankings matter both for a policymaker choosing which corridor to prioritize and for a firm projecting how transit changes will affect a candidate location's long-run profitability. Firms using exogenous-market dynamics will systematically mis-measure candidate locations along the residential-feedback dimension, while firms using endogenous-market dynamics attach value to transit-driven amenity feedback that naive projections cannot see.

That said, a number of limitations in our analyses remain. First, our identification relies on two localized quasi-experiments in a single metropolitan area, and the extent to which our estimates generalize to other, dissimilar environments remains unclear. The bridge collapse provides compelling variation in travel time, but only affected a specific Baltimore-region area; the zoning change provides variation in restaurant supply, but only in another Baltimore area. Second, our model comparison tests whether the endogenous-market specification improves fit relative to the exogenous-market specification, but both specifications share other modeling choices (per-period profit structure, symmetric and homogeneous firms, etc.) that may themselves be misspecified. The model comparison identifies the better-fitting model among those compared, not the correctly specified model in an absolute sense. Third, our household sorting model abstracts away from commuting and other non-amenity determinants of residential location choice; the restaurant amenity and disamenity effects we estimate should be understood as operating conditional on unobserved factors, including workplace proximity, that are absorbed by our location fixed effects. Relatedly, the positive feedback that the model predicts in some disadvantaged areas reflects the direction of the estimated dynamic on average, not a forecast that any particular neighborhood will grow, since the model does not capture the full set of forces that determine long-run neighborhood outcomes.

Future research could extend this framework in several directions. Richer household data would allow estimation of sorting by demographic characteristics, potentially revealing which household types are most responsive to restaurant amenities. Data on restaurant characteristics would enable estimation of quality-differentiated demand, capturing how the type of restaurants (not just the count) affects neighborhood attractiveness. Extending the model to heterogeneous firm types would allow the entry game to capture differences in how chains and independents compete, and could sharpen the counterfactual predictions by allowing policy responses to vary by firm type. Incorporating commuting patterns would allow joint estimation of workplace access and amenity valuation (and restaurant demand), potentially revealing how these forces interact in shaping residential sorting and effective consumer retail markets. Finally, applying this framework to other amenity-providing firms such as retail stores, gyms, or coffee shops would test whether the joint determination documented here is specific to restaurants or reflects a more general phenomenon.

Where businesses locate and where households locate are fundamentally intertwined. A new restaurant does more than serve existing customers; by altering the amenity landscape of its neighborhood, it also reshapes who will live nearby. For managers, our findings suggest that anticipating neighborhood trajectory can be as essential as evaluating current demographics for predicting which locations will support new restaurants. For policymakers, our counterfactual analysis demonstrates that ignoring these feedback dynamics can substantially change the projected firm-surplus response to transit investments, with the direction and magnitude of misprediction depending on local restaurant density and how strongly induced entry feeds back into residential sorting. To build better towns and cities, both for residents and for businesses, it is essential to understand how each part of a neighborhood shapes the others.

References

Adão, Rodrigo, Costas Arkolakis, and Federico Esposito. 2026. "General Equilibrium Effects in Space: Theory and Measurement." *American Economic Journal: Macroeconomics.*

Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential estimation of dynamic discrete games." *Econometrica* 75 (1): 1–53.

Aguirregabiria, Victor, and Junichi Suzuki. 2014. "Identification and counterfactuals in dynamic models of market entry and exit." *Quantitative Marketing and Economics* 12 (3): 267–304.

Ahlfeldt, Gabriel M, Stephen J Redding, Daniel M Sturm, and Nikolaus Wolf. 2015. "The Economics of Density: Evidence From the Berlin Wall." *Econometrica* 83 (6): 2127–2189.

Allen, Treb, and Costas Arkolakis. 2022. "The welfare effects of transportation infrastructure improvements." *The Review of Economic Studies* 89 (6): 2911–2957.

Almagro, Milena, and Tomás Domínguez-Iino. 2025. "Location sorting and endogenous amenities: Evidence from amsterdam." *Econometrica* 93 (3): 1031–1071.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro. 2020. "Transparency in structural research." *Journal of Business & Economic Statistics* 38 (4): 711–722.

Arcidiacono, Peter, Patrick Bayer, Jason R Blevins, and Paul B Ellickson. 2016. "Estimation of dynamic discrete choice models in continuous time with an application to retail competition." *The Review of Economic Studies* 83 (3): 889–931.

Bajari, Patrick, C Lanier Benkard, and Jonathan Levin. 2007. "Estimating dynamic models of imperfect competition." *Econometrica* 75 (5): 1331–1370.

Blevins, Jason R., Ahmed Khwaja, and Nathan Yang. 2018. "Firm Expansion, Size Spillovers, and Market Dominance in Retail Chain Dynamics." *Management Science* 64 (9): 4070–4093.

Bresnahan, Timothy F, and Peter C Reiss. 1991. "Entry and competition in concentrated markets." *Journal of Political Economy* 99 (5): 977–1009.

———. 1994. "Measuring the importance of sunk costs." *Annales d'Économie et de Statistique,* no. 34, 159–180.

Cao, Yue, Judith A Chevalier, Jessie Handbury, Hayden Parsley, and Kevin R Williams. 2024. *Distributional Impacts of the Changing Retail Landscape.* NBER Working Paper 33307. National Bureau of Economic Research.

Caoui, El Hadi, Brett Hollenbeck, and Matthew Osborne. 2025. "Dynamic Entry and Spatial Competition: An Application to Dollar Store Expansion." Forthcoming, *RAND Journal of Economics.*

Chang, Daeyeol, Niyeyesh Meimandi Nejad, Mansoureh Jeihani, and Mansha Swami. 2025. "Beyond the Detour: Modeling Traffic System Shocks After the Francis Scott Key Bridge Failure." *Sustainability* 17 (15): 6916. https://doi.org/10.3390/su17156916.

Couture, Victor, Cécile Gaubert, Jessie Handbury, and Erik Hurst. 2024. "Income Growth and the Distributional Effects of Urban Spatial Sorting." *Review of Economic Studies* 91 (2): 858–898.

Datta, Sumon, and K Sudhir. 2013. "Does reducing spatial differentiation increase product differentiation? Effects of zoning on retail entry and format variety." *Quantitative Marketing and Economics* 11 (1): 83–116.

Davis, Donald R, Jonathan I Dingel, Joan Monras, and Eduardo Morales. 2019. "How segregated is urban consumption?" *Journal of Political Economy* 127 (4): 1684–1738.

Diamond, Rebecca. 2016. "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000." *American economic review* 106 (3): 479–524.

Dingel, Jonathan I, and Felix Tintelnot. 2026. "Spatial Economics for Granular Settings." *Econometrica.*

Donaldson, Dave, and Richard Hornbeck. 2016. "Railroads and American economic growth: A "market access" approach." *The Quarterly Journal of Economics* 131 (2): 799–858.

Fox, Edward J, Steven Postrel, and Amanda McLaughlin. 2007. "The Impact of Retail Location on Retailer Revenues: An Empirical Investigation." Working Paper, Edwin L. Cox School of Business, Southern Methodist University.

Frittelli, John, Ben Goldman, and Ali E. Lohman. 2024. *Baltimore Bridge Collapse: Frequently Asked Questions (FAQ).* CRS Report IF12619. Washington, DC: Congressional Research Service, March. https://www.congress.gov/crs-product/IF12619.

Gechter, Michael, and Nick Tsivanidis. 2025. "Spatial Spillovers from Urban Redevelopments: Evidence from Mumbai's Textile Mills." Working paper.

Gourieroux, Christian, Alain Monfort, and Eric Renault. 1993. "Indirect inference." *Journal of Applied Econometrics* 8 (S1): S85–S118.

Heblich, Stephan, Stephen J Redding, and Daniel M Sturm. 2020. "The making of the modern metropolis: Evidence from London." *The Quarterly Journal of Economics* 135 (4): 2059–2133.

Ifrach, Bar, and Gabriel Y. Weintraub. 2017. "A Framework for Dynamic Oligopoly in Concentrated Industries." *Review of Economic Studies* 84 (3): 1106–1150.

Igami, Mitsuru, and Nathan Yang. 2016. "Unobserved Heterogeneity in Dynamic Games: Cannibalization and Preemptive Entry of Hamburger Chains in Canada." *Quantitative Economics* 7 (2): 483–521.

Luxen, Dennis, and Christian Vetter. 2011. "Real-time routing with OpenStreetMap data." In *Proceedings of the 19th ACM SIGSPATIAL international conference on advances in geographic information systems,* 513–516.

Matthews, John W, and Geoffrey K Turnbull. 2007. "Neighborhood Street Layout and Property Value: The Interaction of Accessibility and Land Use Mix." *Journal of Real Estate Finance and Economics* 35 (2): 111–141.

Miyauchi, Yuhei, Kentaro Nakajima, and Stephen J Redding. 2025. "The economics of spatial mobility: Theory and evidence using smartphone data." *The Quarterly Journal of Economics* 140 (4): 2507–2570.

Monte, Ferdinando, Stephen J Redding, and Esteban Rossi-Hansberg. 2018. "Commuting, Migration, and Local Employment Elasticities." *American Economic Review* 108 (12): 3855–3890.

Newton, Michael A, and Adrian E Raftery. 1994. "Approximate Bayesian inference with the weighted likelihood bootstrap." *Journal of the Royal Statistical Society: Series B (Methodological)* 56 (1): 3–48.

Nishida, Mitsukuni. 2014. "The costs of zoning regulations in retail chains: the case of the City Planning Act of 1968 in Japan." *Journal of Regulatory Economics* 45 (3): 305–328.

———. 2015. "Estimating a model of strategic network choice: the convenience-store industry in Okinawa." *Marketing Science* 34 (1): 20–38.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe. 2023. "What's trending in difference-in-differences? A synthesis of the recent econometrics literature." *Journal of econometrics* 235 (2): 2218–2244.

Rubin, Donald B. 1981. "The Bayesian bootstrap." *The Annals of Statistics* 9 (1): 130–134.

SafeGraph. 2022. *SafeGraph Spend Patterns.* https://doi.org/10.82551/NSF5-R186. https://doi.org/10.82551/NSF5-R186.

Seim, Katja. 2006. "An empirical model of firm entry with endogenous product-type choices." *The RAND Journal of Economics* 37 (3): 619–640.

Shao, Jun, and Dongsheng Tu. 1995. *The Jackknife and Bootstrap.* New York: Springer.

Singh, Vishal P, Karsten T Hansen, and Robert C Blattberg. 2006. "Market Entry and Consumer Behavior: An Investigation of a Wal-Mart Supercenter." *Marketing Science* 25 (5): 457–476.

Suzuki, Junichi. 2013. "Land use regulation as a barrier to entry: Evidence from the Texas lodging industry." *International Economic Review* 54 (2): 495–523.

Tashman, Leonard J. 2000. "Out-of-sample tests of forecasting accuracy: an analysis and review." *International Journal of Forecasting* 16 (4): 437–450. https://doi.org/10.1016/S0169-2070(00)00065-0.

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy. 2008. "Markov Perfect Industry Dynamics with Many Firms." *Econometrica* 76 (6): 1375–1411.

Yang, Hee Jin, Jihoon Song, and Mack Joong Choi. 2016. "Measuring the Externality Effects of Commercial Land Use on Residential Land Value: A Case Study of Seoul." *Sustainability* 8 (5): 432.

A Data Construction

This appendix provides detailed descriptions of data sources and construction procedures summarized in Section 3.

A.1 Summary Statistics

Tables A1–A3 report summary statistics for the main data sources.

Table A1 — Summary Statistics: SafeGraph Spend Panel

	Mean	Median	SD
Monthly spend ($)	5,221	2,488	9,002
Transactions per month	274.8	113	493.4
N (restaurant × month)		63,096	

Notes. Unit of observation is establishment × month. The sample is a balanced panel of 2,629 food-service establishments in Baltimore-Washington metropolitan area observed in all 24 months from January 2023 through December 2024. Data from SafeGraph Spend.

Table A2 — Summary Statistics: Data Axle Restaurant Entry/Exit Panel

Unique establishments	28,850
Firm type	
Independent (1 location)	19,764
Chain (2+ locations)	9,086
Entry and exit (2013–2019)	
Total entries	16,440
Total exits	13,930
CBGs	3,106
Years	2013–2019

Notes. Unit of observation is a unique food-service establishment (NAICS 722) in the 10-county study area active at some point during 2013–2019. Entry is defined as a restaurant's appearance in the database; exit as its disappearance. Firm type is classified by the number of distinct locations sharing the same standardized brand name. Data from Data Axle.

Table A3 — Summary Statistics: CBG Demographics

	Mean	Median
Total population	1,487	1,352
Population density (per sq mi)	7,046	4,520
Median household income ($)	81,420	73,462
Median age	39.6	39.3
% bachelor's degree+ (age 25+)	37.4	34.0
% White	52.9	58.5
N (CBGs)	3,106	

Notes. ACS 2013 5-year estimates (2009–2013) for 3,106 CBGs in the 10-county study area. Population density is computed from Census population and TIGER/Line land area. Education share is among the population aged 25 and over.

A.2 Driving Times and Congestion

We compute driving distances and durations using the Open Source Routing Machine (OSRM), a routing engine built on the OpenStreetMap road network (Luxen and Vetter 2011). Origins are CBG centroids; destinations are restaurant coordinates from the SafeGraph spend panel.

To measure the bridge collapse effect, we compute routing distances on two network configurations. The pre-collapse matrix uses the OpenStreetMap network as of January 2024, with the Francis Scott Key bridge present. The post-collapse matrix uses the same network with the bridge removed. The pre-collapse distance matrix has a mean driving distance of 65.2 km (median 55.5 km, s.d. 49.1 km). The post-collapse matrix is nearly identical in aggregate (mean 65.2 km), but 215,860 OD pairs (1.8% of all pairs) experience route changes. Among affected pairs, the mean increase is 1,096 meters (2.3%); this reflects mechanical re-routing only, with the increased congestion on alternative routes incorporated separately (see Section A.2.1 below).

To define restaurant-level treatment status, we compute each restaurant's minimum driving duration to the bridge location using coordinates for the bridge's north and south anchorages. Restaurants within 30 minutes driving time of the bridge are classified as treated.

A.2.1 Congestion Adjustment

The OSRM routing engine computes shortest-path travel times assuming free-flow conditions. After the bridge collapse, diverted traffic created substantial congestion on alternative routes, increasing travel times beyond what mechanical re-routing implies. We adjust for this using congestion estimates from Chang et al. (2025).

The Chang et al. (2025) estimates are unusually well-suited for our analysis, since they study exactly the event that we exploit, namely the March 26, 2024 collapse of the Francis Scott Key bridge. Using high-resolution probe vehicle data from the Iteris ClearGuide platform, they estimate Travel Time Index (TTI) changes before and after the collapse across three post-collapse periods (the immediate aftermath, fall (September–November 2024), and winter (December 2024–February 2025)) using a set of regression models. Their analysis covers the precise geography and time period of our demand estimation, providing externally validated congestion measures for this event. For their difference-in-differences analysis, they classify corridors within a 10-mile radius of the bridge as "Proximate" (treated) and those beyond 10 miles as "Distant" (control).

We translate these estimates into pair-level congestion multipliers using a two-zone circular geography centered on the bridge.[25] A "metro zone" of radius 20 miles captures the broad area experiencing increased traffic from diverted bridge users, approximating the geographic extent of the 30 corridors studied in that work. A "proximity zone" of radius 10 miles, corresponding to the explicit 10-mile "higher treatment" radius in Chang et al. (2025), captures the additional congestion on areas nearest the bridge site. For each CBG–restaurant pair, we compute the fraction of the straight-line route passing through each zone using geometric line-circle intersection, yielding route fractions f_{gj}^{metro} and f_{gj}^{prox}.[26] The congestion multiplier for pair (g, j) is then:

$$m_{gjt} = 1 + (m_t^{\text{base}} - 1) \cdot f_{gj}^{\text{metro}} + (m_t^{\text{prox}} - m_t^{\text{base}}) \cdot f_{gj}^{\text{prox}} \tag{A.1}$$

where m_t^{base} and m_t^{prox} are time-varying multipliers derived from the Chang et al. (2025) TTI estimates. A route entirely outside both zones receives $m_{gjt} = 1$ (no congestion); a route entirely within the metro zone receives $m_{gjt} = m_t^{\text{base}}$; a route through the proximity zone receives the full proximity multiplier m_t^{prox}.[27]

25. Chang et al. (2025) estimate TTI changes for specific corridors (e.g., I-695, I-95, MD-295), not for arbitrary origin-destination pairs. We approximate their corridor-level estimates with two concentric circular zones to obtain a tractable pair-level adjustment.

26. We use straight-line distances rather than actual routing paths for the zone intersection computation, as the OSRM-computed routes are available only for pre- and post-collapse networks, not for the congestion calculation itself.

27. Specifically, we convert the Chang et al. (2025) regression coefficients to duration multipliers as $m = (\text{TTI}_{\text{pre}} + \Delta\text{TTI})/\text{TTI}_{\text{pre}}$, giving the ratio of post-collapse to pre-collapse travel times. These multipliers vary across the three post-collapse periods, reflecting the gradual dissipation of congestion as drivers adjusted to alternative routes; the largest multipliers occur in the immediate aftermath and decay substantially by fall and winter. (For May–August 2024, we apply the fall-period multiplier.) The adjusted post-collapse duration for pair (g, j) in month t is $\tilde{d}_{gjt} = d_{gj}^{\text{post}} \cdot m_{gjt}$, where d_{gj}^{post} is the OSRM free-flow duration after removing

A.3 Zoning Map Construction

We obtain parcel-level zoning maps from Baltimore County's Comprehensive Zoning Map Process (CZMP) as ESRI shapefiles from the county's open data portal. We focus on zoning classifications that permit food-service establishments. Between the 2012 and 2016 maps, 386 parcels were newly classified into restaurant-permitting zones.

We construct a CBG-level measure of nearby restaurant-permitting zoning using an annulus approach. For each CBG, we compute the share of land area within 0–1-mile, 1–2-mile, and 2–3-mile annuli around the CBG centroid that is zoned for restaurant use, separately for the 2012 and 2016 maps. The change in restaurant-zoned share between 2016 and 2012 within each annulus is

$$\Delta Z_g^{a\text{-}b} = \frac{\mathrm{Area}(A_g^{a,b} \cap \mathcal{Z}_{2016})}{\mathrm{Area}(A_g^{a,b})} - \frac{\mathrm{Area}(A_g^{a,b} \cap \mathcal{Z}_{2012})}{\mathrm{Area}(A_g^{a,b})} \tag{A.2}$$

where $A_g^{a,b}$ is the annulus between a and b miles from CBG g's centroid and $\mathcal{Z}_t$ is the union of restaurant-permitting zone polygons in year t.

We observe 502 CBGs in Baltimore County across each of the three distance rings. The mean restaurant-zoned share within the 1-mile ring is 10.4% in 2012 and 10.5% in 2016; the mean change is 0.037 percentage points (s.d. 0.23 percentage points). For details on the CZMP process and identification strategy, see Appendix B.2.1.

A.4 American Community Survey

We obtain CBG-level demographic characteristics from the American Community Survey (ACS) 5-year estimates. For demand estimation, we use the 2024 vintage (covering survey years 2020–2024) for Maryland and Washington, DC. For the zoning analysis, we use annual ACS vintages from 2013 through 2023 to construct a CBG-level population panel.

the bridge. For distant restaurants whose straight-line routes do not intersect either zone, the route fractions are zero and the multiplier reduces to 1.0.

We use total population to construct the CBG-level population panel for the household location choice model, and median household income, educational attainment (percent with a bachelor's degree or higher among population aged 25 and over), and population density as CBG-level demographic controls.

B Identification and Reduced-Form Robustness

This appendix provides further discussion and background on the reduced-form identification strategies, and presents a series of robustness tests for the two quasi-experiments used in the main text.

B.1 Further Discussion on Bridge Collapse Identification

The causal interpretation of the two-way fixed effects estimates in Section 4 rests on the standard assumptions of parallel trends, no anticipation, and the stable unit treatment value assumption, or SUTVA (Roth et al. 2023).

First, Figure 2 provides direct evidence that treated restaurants (within 30 minutes of the bridge) and control restaurants (30–60 minutes) trended in parallel prior to the collapse, which suggests that they would have trended in parallel absent the collapse, as pre-period coefficients cluster tightly around zero with no systematic pre-trend across the nine months preceding the collapse in both the full sample (Panel A) and the intermediate-annulus-excluded sample (Panel B). The same evidence supports the "no anticipation" assumption; further, anecdotally, the cause of the collapse (the container ship *Dali* lost power without warning at 1:28 a.m. on March 26, 2024, striking the Francis Scott Key Bridge and causing its immediate collapse) was genuinely not anticipated by any of the actors in the area.

However, as noted in the main text, we acknowledge that the reduced-form estimates may be affected by SUTVA violations. First, control restaurants 30–60 minutes from the bridge may themselves have experienced travel-time increases from collapse-induced congestion (Chang et al. 2025); this would attenuate the estimated treatment effect, by shrinking the post-period comparative gap between treatment and control revenue. Consistent with mild attenuation, the intermediate annulus exclusion in Table 1 (columns 2 and 4) yields slightly larger point estimates when the 30–45 minute zone is dropped from the control group. Second, demand substitution from treated to control restaurants could inflate the DiD if spending redirected to control-group restaurants raises their revenue. As noted in the main text, these potential violations motivate our use of the structural model, which can capture and appropriately account for both of these patterns, while still leveraging the quasi-experimental variation from the bridge collapse to identify effects. The model also allows us to account for congestion effects, leveraging the estimates of Chang et al. (2025).

We also point out that the DC placebo (Appendix B.10, Table B8) provides evidence that neither channel is significantly large; the placebo specification uses Baltimore-area restaurants beyond 30 minutes as controls and so would pick up any significant movement in the control group. If congestion depressed control-group spending, the placebo DiD would be significantly positive (DC flat, controls down), while if demand substitution elevated control-group spending, the placebo DiD would be significantly negative (DC flat, controls up). The null placebo estimates are consistent with neither channel operating at meaningful magnitudes in this comparison.

B.2 Further Discussion on Zoning Reform Identification

B.2.1 Rezoning Institutional Details

Baltimore County's Comprehensive Zoning Map Process (CZMP) is a four-year county-wide rezoning cycle codified in Baltimore County law. Identification rests on two statutory features of the CZMP. The four-year cycle timing is set by county code and does not respond to local market conditions, and between-cycle rezoning faces substantially higher legal barriers and automatic reversion provisions, making the CZMP the dominant source of zoning variation.[28]

As discussed in the main text, which specific parcels are rezoned is almost certainly non-random, as individual petitions may reflect local development pressures or property-specific considerations. The exogenous policy timing does, however, suggest that the event study pre-trend test is a reliable diagnostic for confounding, since any omitted variable that correlates

28. Between cycles, property owners may petition through a separate process, but must demonstrate a "substantial change in the character of the neighborhood" or a "mapping error." Approved changes revert if not developed within three years.

with rezoning location would need to also align precisely with the predetermined four-year schedule to generate spurious post-period effects without also appearing in the pre-period. The flat pre-trends we document in Figure 3 are consistent with no such alignment. The robustness tests that follow further probe this by examining whether the results are driven by specific geographic areas where confounding development activity is most plausible. We further note that our downstream household model separately includes non-restaurant retail density as an auxiliary control (Section 5), so any residual generic commercial-development effects that survive these diagnostics are absorbed downstream rather than mischarged to restaurant proximity.

B.2.2 Identifying Assumptions

As with the Francis Scott Key bridge collapse, the causal interpretation of the two-way fixed effects estimates in Section 5 rests on the standard assumptions of parallel trends, no anticipation, and SUTVA.

For parallel trends, as discussed in the prior subsection and tested further in the robustness checks below, we argue that the null pre-trends in Figure 3 provides credible evidence of parallel trends, as pre-period coefficients for 2013 and 2014 are flat and tightly estimated across all three distance rings. The same evidence supports no anticipation, as we do not see any effects on restaurants or population until the (exogenous, policy-determined) timing of the rezoning, and population effects either dovetail with or shortly follow after restaurant density changes. The estimates are also robust to a wide battery of robustness checks, detailed below.

Regarding SUTVA violations, one may be concerned that areas farther away lose (gain) population when residents move into (away from) more or less attractive areas, respectively, due to changed area amenity density; if large, this could spuriously inflate reduced-form effect estimates by increasing the post-period difference between treated and control areas. In response to this, we first note that the econometric specification is already set up to account for effects across wide radii, controlling for effects of zoning changes 2-3 miles away when estimating effects for zoning changes 0-1 or 1-2 miles away; and further, the pattern of empirical results suggests that zoning changes do not significantly affect population density in the 2-3 mile radius, suggesting that any effects on control areas are likely dispersed and not significantly large. Second, while we acknowledge that our reduced-form estimates may nonetheless suffer from some degree of bias in this way, we argue (as with our distance-sensitive demand model above) that this is only greater motivation for our use of a structural model, estimated with indirect inference from this same source of identifying variation, which can directly account for any possible effects on "control" units when reproducing the reduced-form effect and backing out the relevant preference parameters over area amenities.

B.3 Corridor Exclusion

We identify three corridors with active commercial investment during the CZMP cycle (Towson/Timonium in the north, Catonsville/Arbutus in the southwest, and Dundalk/Essex in the southeast; see Figure B1) and re-estimate the baseline specification after excluding each one individually. Table B1 reports the results.

Table B1 — Corridor Exclusion Robustness Test

	Baseline (1)	Excl. Towson (2)	Excl. Catonsville (3)	Excl. Dundalk (4)	Excl. All Three (5)
0–1 mi × Post	$-5.31^{\dagger}$	$-5.96^{\dagger}$	-5.48^{*}	-6.04	-9.47
	(2.78)	(3.50)	(2.79)	(4.01)	(6.42)
1–2 mi × Post	10.19^{*}	$11.05^{\dagger}$	9.60^{*}	$14.40^{\dagger}$	$23.31^{\dagger}$
	(4.65)	(5.72)	(4.62)	(7.37)	(12.71)
CBG FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
CBGs excluded	0	119	63	111	293

Notes. Dependent variable is log population. Column (1) reproduces the baseline specification from Table 3. Columns (2)–(4) each exclude one development corridor; column (5) excludes all three simultaneously. All specifications include all three annuli jointly; only the 0–1 and 1–2 mile coefficients are reported. Standard errors clustered at CBG level in parentheses. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Figure B1 — Corridor Exclusion: Development Corridors and Exclusion Maps

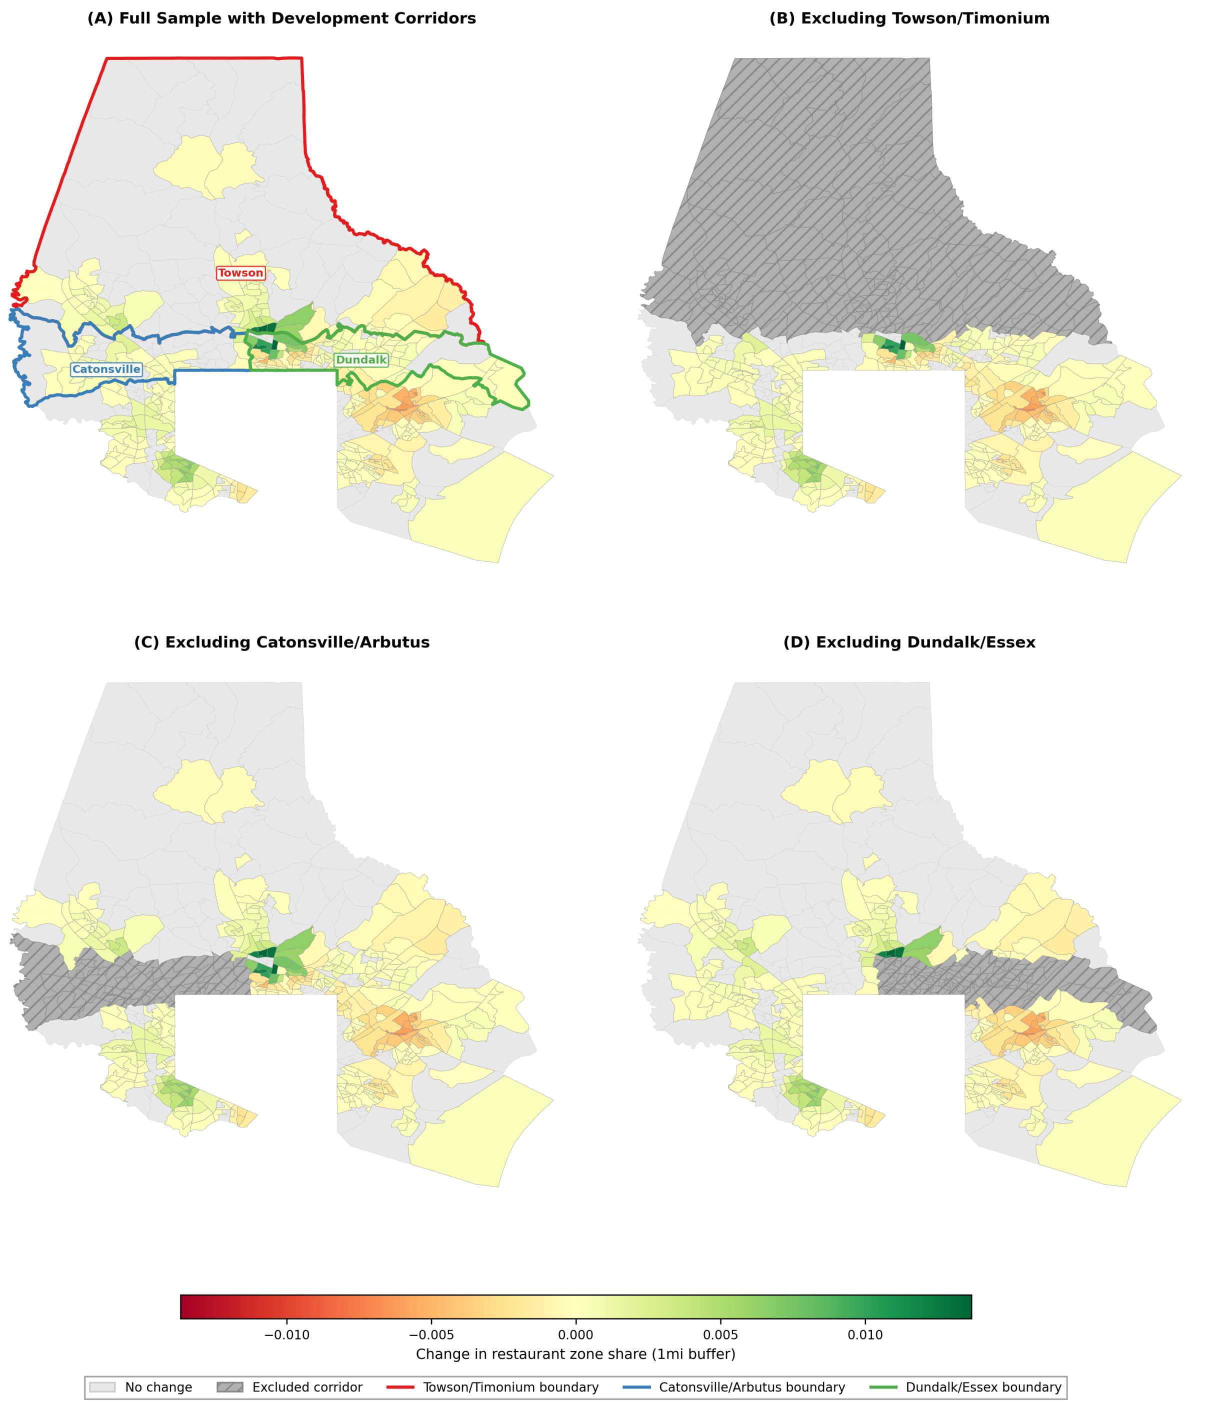

Notes. Panel (A) shows the full sample with three development corridors outlined. Panels (B)–(D) each exclude one corridor (hatched). Treatment is the change in restaurant-zoned land share within a 1-mile buffer of each CBG between 2012 and 2016. Color scale: green indicates gained restaurant zoning, red indicates lost restaurant zoning, gray indicates no change.

The sign pattern is preserved in every specification: the 0–1 mile coefficient remains negative and the 1–2 mile coefficient remains positive regardless of which corridor is dropped. Point estimates are stable or, if anything, larger in magnitude when corridors are excluded, though standard errors widen given the reduced number of treated CBGs. Column (5) excludes all three corridors simultaneously; even in this restricted sample, the non-monotonic pattern persists with point estimates roughly twice the baseline magnitude and the 1–2 mile amenity remaining marginally significant ($p = 0.07$). No single corridor or combination of corridors drives the results, confirming that the baseline estimates are not driven by concurrent neighborhood investment.[29]

B.4 Treatment Cluster Exclusion

While the corridor exclusion above uses broad geographic definitions, we can also perform a more targeted test using data-driven clusters. We use DBSCAN spatial clustering to identify three tight geographic clusters of high-treatment-intensity CBGs (those with $|\Delta\text{zone share}| > 0.002$): a Towson cluster in the north (63 CBGs), a Catonsville cluster in the southwest (24 CBGs), and a Dundalk cluster in the southeast (63 CBGs). These clusters, shown in Figure B2, correspond to the areas with the most visually prominent treatment on the zoning change map. We then re-estimate the baseline specification after excluding each cluster individually.

Table B2 reports the results. The sign pattern is preserved in every specification: the 0–1 mile coefficient remains negative and the 1–2 mile coefficient remains positive regardless of which cluster is dropped. Magnitudes are stable, with the Towson cluster contributing the most identifying variation (standard errors roughly double when it is excluded). The Catonsville exclusion, despite dropping only 24 CBGs, barely moves the point estimates,

29. The non-monotonic spatial pattern is also difficult to explain through generic commercial development, which would plausibly generate monotonically declining effects with distance; separately, our downstream household model includes non-restaurant retail density as an auxiliary control, absorbing any generic commercial-density effect on population.

Figure B2 — Leave-One-Cluster-Out: Treatment Clusters and Exclusion Maps

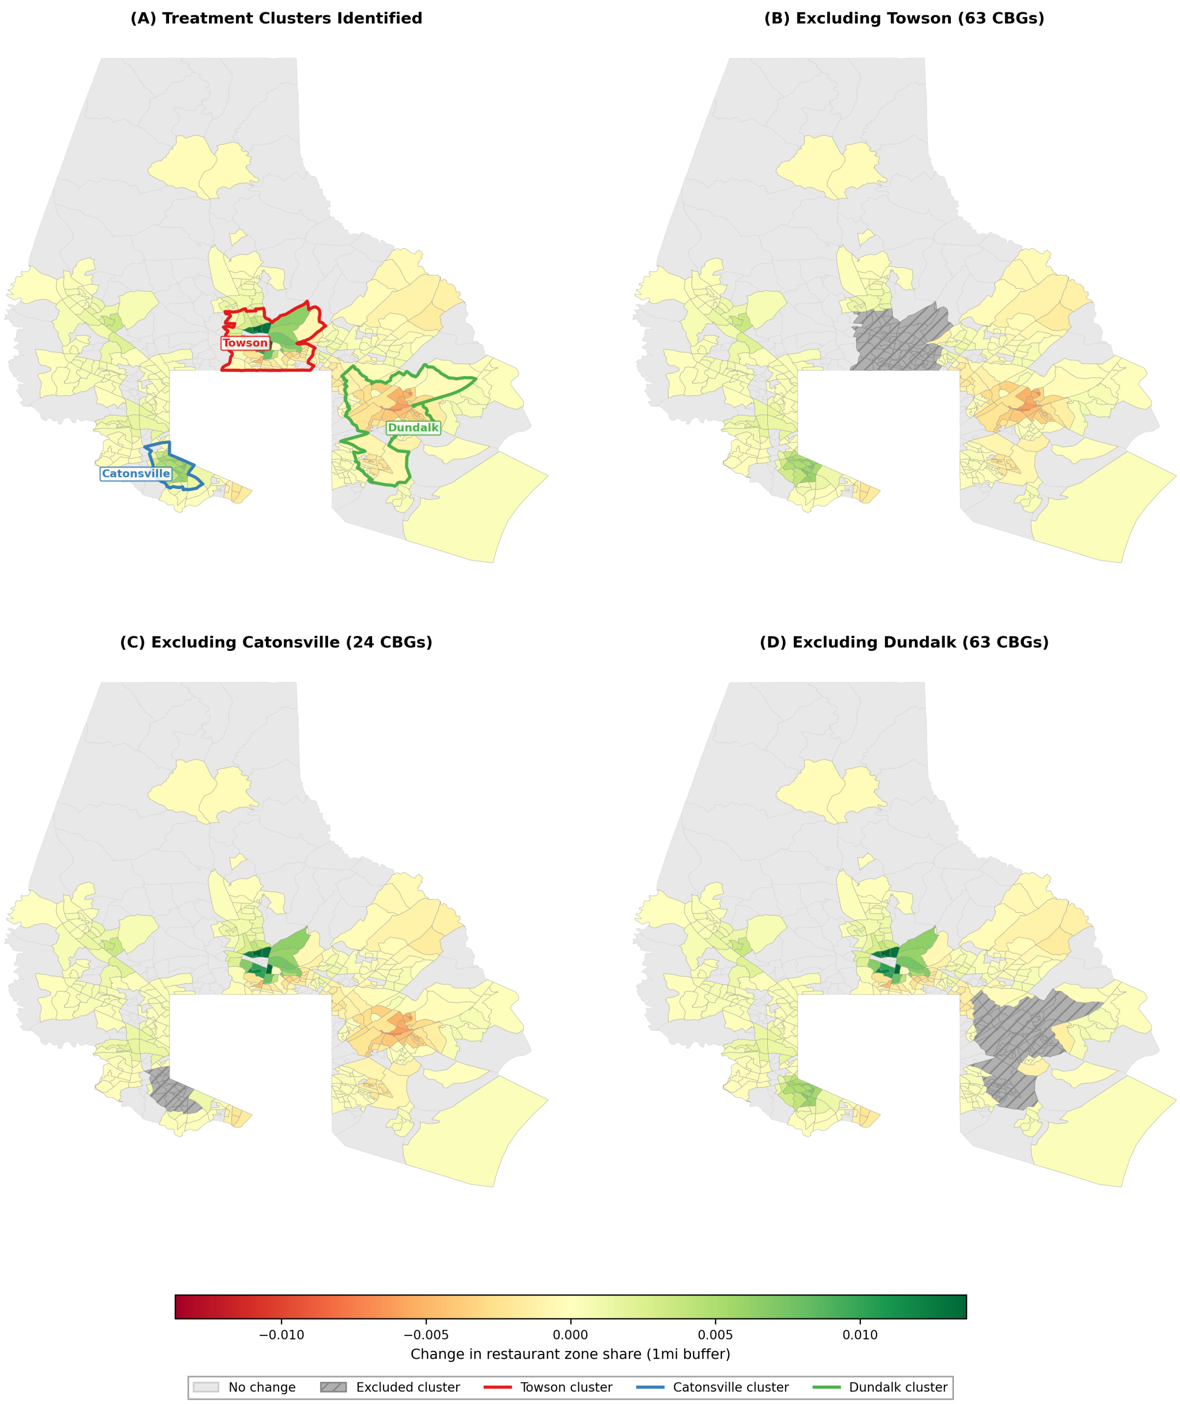

Notes. Panel (A) shows the three treatment clusters identified via DBSCAN on CBGs with $|\Delta\text{zone share}| > 0.002$. Panels (B)–(D) each exclude one cluster (hatched). Treatment is the change in restaurant-zoned land share within a 1-mile buffer of each CBG between 2012 and 2016. Color scale: green indicates gained restaurant zoning, red indicates lost restaurant zoning, gray indicates no change.

Table B2 — Treatment Cluster Exclusion Robustness

	Baseline	Excl. Towson Cluster	Excl. Catonsville Cluster	Excl. Dundalk Cluster
0–1 mi × Post	$-5.31^{\dagger}$	-7.29	-4.37	$-5.07^{\dagger}$
	(2.78)	(5.73)	(2.67)	(2.95)
1–2 mi × Post	10.19^{*}	12.25	$8.83^{\dagger}$	7.14
	(4.65)	(9.58)	(4.74)	(4.48)
CBG FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
CBGs excluded	0	63	24	63

Notes. Dependent variable is log population. Column (1) reproduces the baseline specification. Columns (2)–(4) each exclude one treatment cluster identified via DBSCAN spatial clustering on high-treatment-intensity CBGs ($|\Delta\text{zone share}| > 0.002$). All specifications include three annuli jointly; only the 0–1 and 1–2 mile coefficients are reported. Standard errors clustered at CBG level in parentheses. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

while excluding Dundalk modestly attenuates the 1–2 mile effect. No single cluster drives the non-monotonic pattern.

B.5 Ring Boundary Sensitivity

The preceding tests address potential geographic confounds; we now consider whether the non-monotonic pattern may be sensitive to the particular 1-mile cutoff used to define the distance annuli. Table B3 addresses this by perturbing the inner boundary by ±0.1 miles and the outer boundary of the middle ring by ±0.1 miles. In all cases, we recompute buffer-polygon intersections using the 2012 and 2016 Baltimore County zoning maps and estimate the same two-way fixed effects DiD specification. The qualitative pattern (negative near, positive middle) is stable across all specifications. Shifting the inner boundary to 0.9 or 1.1 miles preserves the negative near-ring coefficient and the significant positive middle-ring coefficient. Similarly, shifting the outer boundary to 1.9 or 2.1 miles leaves the middle-ring amenity effect significant at the 5% level.

Table B3 — Ring Boundary Sensitivity: Zoning Effects on Population

	Distance Annulus	
Ring boundaries	Near	Middle
Baseline		
0–1 / 1–2 mi	$-5.31^{\dagger}$	10.19^{*}
	(2.78)	(4.65)
Perturbing inner boundary		
0–0.9 / 0.9–2 mi	$-4.55^{\dagger}$	9.28^{*}
	(2.45)	(4.73)
0–1.1 / 1.1–2 mi	-4.25	9.09^{*}
	(2.91)	(4.55)
Perturbing outer boundary		
0–1 / 1–1.9 mi	$-5.31^{\dagger}$	8.89^{*}
	(2.78)	(4.17)
0–1 / 1–2.1 mi	$-5.19^{\dagger}$	11.67^{*}
	(2.76)	(5.00)
CBG FE	Yes	
Time FE	Yes	

Notes. Two-period DiD estimates as in Table 3, perturbing the baseline 1-mile inner boundary and 2-mile outer boundary by ±0.1 miles. Treatment is the change in restaurant-zoned land share within each annulus between 2012 and 2016. All specifications include the full set of annuli jointly; only the 0–1 and 1–2 mile coefficients are reported. Standard errors in parentheses, clustered at the CBG level. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

B.6 Inner Ring Exclusion

A natural concern with the 0–1 mile disamenity finding is that it could be driven by restaurants literally displacing residents from the same parcel or immediate block, rather than reflecting a broader neighborhood-level disamenity from noise and congestion. We address this with two tests that progressively hollow out the innermost portion of the 1-mile treatment ring. In the first, we subtract each CBG's own zoning overlap from its 1-mile buffer, so that the treatment captures only restaurant zoning in neighboring CBGs. In the second, we exclude the entire 0–0.25 mile ring, constructing a 0.25–1 mile annulus that removes all immediately adjacent zoning changes.

Table B4 — Inner Ring Exclusion: Zoning Effects on Population

	Baseline (1)	Excl. own-CBG (2)	0.25–1 mi annulus (3)
0–1 mi × Post	$-5.31^{\dagger}$	$-4.70^{\dagger}$	$-5.13^{\dagger}$
	(2.78)	(2.62)	(2.69)
1–2 mi × Post	10.19^{*}	10.01^{*}	10.14^{*}
	(4.65)	(4.63)	(4.65)
CBG FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes

Notes. Two-period DiD estimates as in Table 3. Column (1) reproduces the baseline specification. Column (2) subtracts each CBG's own zoning overlap from its 1-mile buffer, so that the 0–1 mile treatment reflects only neighboring CBGs' zoning. Column (3) excludes the inner 0.25 miles entirely, using a 0.25–1 mile annulus. All specifications include all three annuli jointly; only the 0–1 and 1–2 mile coefficients are reported. Standard errors in parentheses, clustered at the CBG level. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Table B4 reports the results. Excluding own-CBG zoning attenuates the coefficient by 12% (-4.70 versus -5.31), confirming that the disamenity is driven primarily by restaurants in the surrounding neighborhood rather than on the same block. The 0.25–1 mile annulus specification yields a coefficient close to baseline (-5.13, $p = 0.057$, versus baseline -5.31), indicating that the innermost quarter-mile contributes little treatment variation and that the disamenity operates at the neighborhood scale. The 1–2 mile coefficients are virtually unchanged across all three specifications.

B.7 Baltimore City Boundary Exclusion

Baltimore County wraps around Baltimore City, so one may also wonder whether treated CBGs near the county–city boundary reflect city-specific trends rather than the zoning treatment. Table B5 addresses this concern by progressively excluding Baltimore County CBGs whose centroids lie within 0.25 and 0.5 miles of the city boundary. The 0–1 mile disamenity, which is significant only at the 10% level in the baseline (-5.31, $p = 0.056$), strengthens to 5% significance under both exclusions (-6.30, $p = 0.024$ at 0.25 miles; -5.43, $p = 0.049$ at 0.5 miles). The 1–2 mile amenity remains stable at 9.2 to 10.6 across

Table B5 — Baltimore City Boundary Exclusion: Zoning Effects on Population

	Baseline (1)	Excl. 0.25 mi (2)	Excl. 0.5 mi (3)
0–1 mi × Post	$-5.31^{\dagger}$	-6.30^{*}	-5.43^{*}
	(2.78)	(2.78)	(2.76)
1–2 mi × Post	10.19^{*}	10.58^{*}	$9.17^{\dagger}$
	(4.65)	(4.76)	(4.89)
CBG FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
CBGs	3,094	3,066	3,032

Notes. Two-period DiD estimates as in Table 3. Each column excludes Baltimore County CBGs whose centroids lie within the specified distance of the Baltimore City boundary; control-county CBGs are retained regardless of proximity. Both specifications include all three annuli jointly; only the 0–1 and 1–2 mile coefficients are reported. Standard errors clustered at the CBG level in parentheses. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

all specifications. If anything, removing city-adjacent CBGs sharpens identification by eliminating noisy observations, suggesting that the baseline estimates are not driven by proximity to Baltimore City.

B.8 Rent Robustness

A natural concern is that the population shifts documented above reflect rent-driven displacement rather than amenity-based sorting. If rezoning raises nearby rents, households may be pushed out of proximate CBGs rather than voluntarily relocating in response to restaurant amenities. To assess this, we re-estimate the same two-period DiD specification from equation (7) using log median rent (ACS variable B25064) as the outcome. If rent-driven displacement were the primary mechanism, we would expect rent increases in the 0–1 mile ring (pushing residents out) and decreases or null effects at 1–2 miles (pulling residents in), the opposite sign pattern from the population results.

Table B6 reports the results. The 0–1 mile coefficient is negative (-5.65, $p = 0.074$), indicating that rents, if anything, decline near rezoned areas. This is the same direction as the population effect, ruling out a rent-increase displacement story, since residents are

Table B6 — Rent Robustness: Effect of Restaurant Zoning on Log Median Rent

	Log Median Rent		
	0–1 mi	1–2 mi	2–3 mi
Effect of zoning	$-5.65^{\dagger}$	4.08	$-13.91^{\dagger}$
	(3.16)	(4.41)	(7.22)
CBG FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
CBGs	2,364	2,364	2,364

Notes. Two-period DiD estimates as in Table 3, with log median rent (ACS B25064) replacing log population as the outcome. CBG-years with missing rent values are dropped. Standard errors clustered at the CBG level in parentheses. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

not being priced out of nearby CBGs. The 2–3 mile coefficient is also negative (-13.91, $p = 0.054$), while the 1–2 mile coefficient is positive but statistically insignificant (4.08, $p = 0.354$). These patterns are inconsistent with rent-driven displacement and instead suggest that the population responses documented in Table 3 reflect amenity-based sorting.

B.9 Covariate Balance

A natural question is whether omitting area-level covariates such as income and education from the household sorting model biases the restaurant proximity estimates. As discussed in Section 5, our identifying moments are derived from a DiD specification with CBG and year fixed effects, which absorbs time-invariant location differences. A remaining concern is that the zoning reform may have differentially shifted neighborhood demographics in ways correlated with restaurant density changes, creating omitted variable bias.

We test this directly by re-estimating the same DiD specification from equation (7) using log median household income (ACS B19013) and education share (fraction with BA or higher, ACS B15003) as outcomes. Results are presented in Table B7. No coefficient is statistically significant across either outcome or any distance ring, confirming that the zoning reform does not move area demographics. The omission of these covariates from the structural model is accordingly unlikely to confound the restaurant proximity estimates.

Table B7 — Covariate Balance: Effect of Restaurant Zoning on Income and Education

	Log Median Income	Education Share (BA+)
Zoning × Post (0–1 mi)	3.00	1.64
	(3.19)	(1.52)
Zoning × Post (1–2 mi)	6.64	−2.30
	(4.63)	(2.40)
Zoning × Post (2–3 mi)	3.10	0.93
	(7.36)	(2.72)
CBGs	3,079	3,091
CBG FE	Yes	Yes
Year FE	Yes	Yes

Notes. Two-period DiD estimates as in Table 3, with log median household income (ACS B19013) and education share (fraction BA+ from ACS B15003) replacing log population as outcomes. Standard errors clustered at the CBG level in parentheses.

B.10 Other-City Placebo Test

The robustness tests above address potential confounds in the zoning identification. We now turn to the consumer demand identification, where a potential concern is that the spending decline reflects a broader regional or seasonal trend rather than the bridge collapse specifically. To rule this out, we conduct a placebo test using restaurants in Washington, DC, which are located roughly 40 miles from the Francis Scott Key Bridge and should be entirely unaffected by its collapse. If the same specification produced a significant effect in DC, it would suggest a problem with the identification strategy.

The placebo reverses the treatment assignment: DC restaurants that are far from the bridge and unambiguously unaffected by its collapse become the "treated" group, while Baltimore-area restaurants that are moderately far serve as controls. Specifically, we designate all 1,475 DC restaurants as placebo-treated and use Baltimore-area restaurants more than 30 minutes from the bridge as controls. Because the DC and Baltimore samples are drawn from overlapping geographic bounding boxes, 180 restaurants appear in both datasets; we exclude these from the control group to avoid contamination, leaving 1,301 controls. We then

Figure B3 — Other-City Placebo: DC Restaurant Spending Event Study

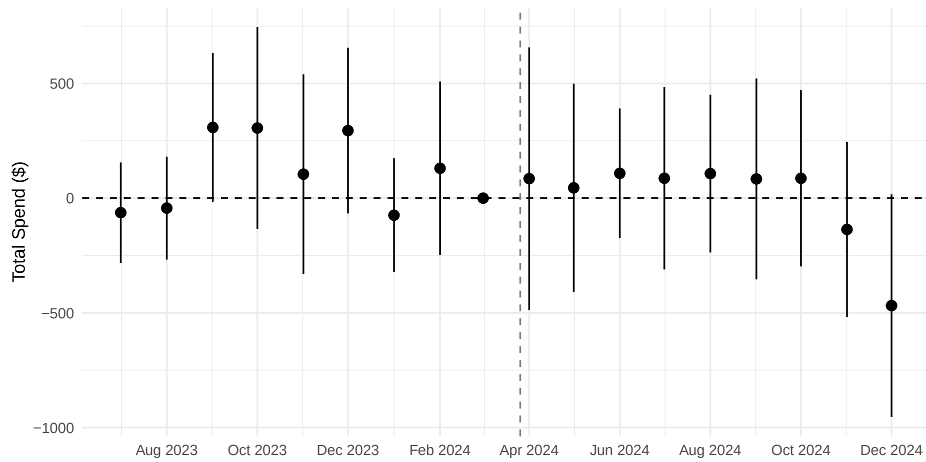

Notes. Monthly event study coefficients for total restaurant spending, using DC restaurants as "treated" and Baltimore restaurants more than 30 minutes from the bridge as controls. Vertical dashed line indicates the bridge collapse (March 26, 2024). Bars show 95% confidence intervals with restaurant-clustered standard errors. No individual coefficient is significant at the 5% level.

estimate the same two-way fixed effects DiD specification on total monthly spending, with restaurant and month fixed effects and standard errors clustered at the restaurant level.

Table B8 reports the static DiD estimates from a symmetric three-month window around the bridge collapse, matching column (1) of Table 1. The main Baltimore specification yields a significant decline, while the DC placebo yields null effects. Figure B3 presents the event study over the full nine-month window; no individual month is statistically significant at the 5% level, and the pre-period coefficients show no systematic pattern, supporting the validity of the parallel trends assumption in the main analysis.

B.11 Bridge-Proximity Exclusion

A potential concern with the consumer demand identification is that the spending decline may partly reflect first-order disruption near the bridge site (construction activity or road closures) rather than the network-wide increase in travel times that our structural model captures. If this were the case, the DiD estimate would overstate the causal effect of travel time on restaurant spending.

To address this, we exclude restaurants that are geographically close to the bridge, where

Table B8 — Other-City Placebo: Baltimore vs. DC

	Baltimore (1)	DC Placebo (2)
Treat × Post	-223.70^{**}	60.83
	(83.78)	(132.41)
Restaurant FE	Yes	Yes
Time FE	Yes	Yes
Observations	15,774	16,656

Notes. Two-way fixed effects DiD estimates of the bridge collapse on total monthly restaurant spending, using a symmetric three-month window (January–June 2024). Column (1) reproduces the Baltimore estimate from Table 1, column (1). Column (2) designates DC restaurants as "treated". Controls are Baltimore restaurants more than 30 minutes from the bridge; the control count is lower in column (2) because 180 DC-area restaurants that also appear in the Baltimore sample are excluded to avoid the same restaurant appearing as both treated and control. Standard errors in parentheses, clustered at the restaurant level. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

such first-order disruption effects would be most concentrated. For each restaurant, we compute straight-line (haversine) distance from the Francis Scott Key Bridge and re-estimate the baseline binary DiD after excluding all restaurants within 5 miles. The control group (restaurants more than 30 minutes driving time from the bridge) is unaffected by this exclusion, as all control restaurants are well beyond 5 miles of the bridge.

Table B9 reports the results. Excluding the 121 treated restaurants within 5 miles of the bridge (those most plausibly exposed to direct bridge-site disruption) actually strengthens the estimate slightly ($-\$239$, $p = 0.014$, versus baseline $-\$219$). This is consistent with the spending decline propagating through the road network via increased travel times rather than reflecting localized disruption near the bridge.

Table B9 — Bridge-Proximity Exclusion

	Baseline	Excl. <5 mi
Treat × Post	-218.69^{*}	-238.76^{*}
	(94.55)	(96.62)
Restaurant FE	Yes	Yes
Time FE	Yes	Yes
Observations	47,322	45,144

Notes. Dependent variable is monthly total spending ($). The second column re-estimates the baseline binary DiD from Table 1 after excluding treated restaurants within 5 miles straight-line (haversine) distance of the Francis Scott Key Bridge. The control group is unchanged, as all control restaurants exceed 5 miles from the bridge. All specifications use the 9-month window (July 2023–December 2024) with intermediate annulus included. Standard errors in parentheses, clustered at the restaurant level. $^{\dagger}p < 0.10$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

C Observational Estimate Comparisons

As discussed in the main text, our quasi-experimental identification yields different estimates than observational approaches, particularly for consumer demand. This appendix presents the full comparison for both models.

C.1 Consumer Demand Model

Table C1 compares the quasi-experimental and orthogonality-based estimates of the consumer demand model. Both specifications use the same sample of restaurants and impose a constant α. The quasi-experimental specification identifies β^d by matching the model-implied DiD to the empirical estimate from the bridge collapse, using a 9-month symmetric window. The orthogonality specification replaces the DiD moment with a cross-sectional condition requiring that spending prediction errors be uncorrelated with log travel duration (see Appendix E for the moment conditions). Unlike the DiD moment, this condition does not exploit variation from the bridge collapse, and instead identifies β^d purely from the cross-sectional covariation between travel duration and demand residuals.

The orthogonality estimate is $\hat{\beta}^d = 2.08$, which has the opposite sign from the DiD-identified estimate (-0.39). Without the within-restaurant variation provided by the bridge collapse, the estimator recovers the equilibrium cross-sectional relationship between travel duration and demand rather than the causal effect of travel cost. In the cross-section, higher-quality restaurants tend to draw customers from a wider geographic catchment, and consumers willingly travel farther for better-matched restaurants, generating a positive correlation between unobserved restaurant quality and average travel duration. Restaurants may also tend to locate nearest to their most marginal customers rather than to their most loyal ones, reinforcing this positive correlation. The orthogonality estimator, unable to separate this equilibrium sorting from the causal effect of distance, produces a positive $\hat{\beta}^d$. The DiD-identified estimate avoids this confound by using within-restaurant variation in travel time induced by the bridge collapse, which is orthogonal to time-invariant restaurant quality.

Table C1 — Consumer Demand Model: DiD-Identified vs. Observational

	DiD-Identified	Observational
Travel Time (β^d)	-0.39	2.08
Moment	DiD match	$\mathbb{E}[r \times \log d] = 0$
N restaurants	2,629	2,629
N observations	47,322	47,322

Notes. Estimates from logit demand model with constant α. Column 1 identifies β^d by matching the model-implied DiD coefficient to the empirical bridge collapse estimate (indirect inference) using a 9-month symmetric window with the intermediate annulus included. Column 2 replaces the DiD moment with a cross-sectional orthogonality condition requiring that average spending prediction errors be uncorrelated with log travel duration; this condition does not exploit bridge collapse variation. Both columns use the same sample. Travel time enters in logs.

C.2 Household Location Choice Model

Table C2 compares the quasi-experimental and observational estimates of the household location choice model. The quasi-experimental specification matches indirect-inference DiD moments from the zoning change. The observational specification omits the DiD moments, relying only on the mobility rate, same-county rate, and county flow moments for identification.

Without the quasi-experimental variation, the parameter estimates change substantially. The moving cost grows in magnitude from -6.71 to -10.10, the 0–1 mile disamenity grows from -0.038 to -0.252, the 1–2 mile amenity shrinks from 0.023 to 0.012, and the distance preference weakens from -4.25 to -2.48. Normalized against the moving cost, the 0–1 mile disamenity grows by roughly 4.4-fold while the 1–2 mile amenity shrinks by roughly 2.9-fold. This is consistent with the observational moments capturing equilibrium correlations between population and restaurant density rather than the causal effect that the zoning shock isolates: without variation that pins down the response of households to local restaurant supply, the model fits residual patterns in the population panel that reflect both true preferences and long-run sorting.

Table C2 — Household Location Choice Model: DiD-Identified vs. Observational

Parameter	Zoning DiD Moments Included	No DiD Moments Included
Moving cost (β^{move})	-6.71	-10.10
Restaurant proximity, 1 mi (β_1^{rest})	-0.038	-0.252
Restaurant proximity, 2 mi (β_2^{rest})	0.023	0.012
Distance preference (β^{dist})	-4.25	-2.48
N CBGs	3,106	3,106

Notes. Estimates from GMM matching simulated moments to empirical targets, both using the baseline household specification (r5_ldrf: log-distance driving-time metric, same-county fixed effect, non-restaurant density in both annuli, log-density/log-income push–pull terms, distance-to-CBD). Quasi-experimental column matches DiD moments from the zoning change. Observational column omits the DiD moments, relying only on the mobility rate, same-county rate, and county flow moments for identification. Restaurant proximity coefficients measure the effect of restaurant count within each distance ring on CBG utility. Distance preference is per 10 miles.

C.3 Cross-Period Stability of Observational Preference Estimates

Our analysis combines preference estimates identified in different time periods: the consumer demand model uses 2023–2024 spending data (exploiting the bridge collapse), while the household sorting and entry models use 2013–2019 data (exploiting the zoning change). This appendix provides evidence that the relevant preferences are broadly stable across these periods.

C.3.1 Household Sorting Model

We re-estimate the household sorting model on 2022–2023 ACS data using the observational specification (omitting DiD moments, which are unavailable outside the zoning window). Because absolute parameter magnitudes depend on normalization, we compare the ratio of each preference coefficient to the moving cost across periods; these ratios capture the economically meaningful tradeoffs that determine residential sorting.

Table C3 reports the results. The 0–1 mile disamenity ratio and distance sensitivity ratio are close in magnitude across periods (0.025 vs. 0.029 and 0.246 vs. 0.269, respectively); the 1–2 mile amenity ratio is essentially zero in both periods, consistent with the fact that this coefficient is only sharply identified by the zoning DiD moments that the observational specification omits. We also compare reduced-form cross-sectional regressions of log population on log restaurant counts and CBG characteristics across periods; none of the coefficient differences are statistically significant (all $p > 0.19$).

Table C3 — Cross-Period Stability of Household Preference Ratios

	2013–2019	2022–2023
$\lvert\beta_1^{\text{rest}}\rvert / \lvert\beta^{\text{move}}\rvert$	0.025	0.029
$\beta_2^{\text{rest}} / \lvert\beta^{\text{move}}\rvert$	0.001	-0.005
$\lvert\beta^{\text{dist}}\rvert / \lvert\beta^{\text{move}}\rvert$	0.246	0.269

Notes. Each ratio divides the indicated coefficient by the absolute value of the moving cost from the same estimation period. Both periods use the observational specification (no DiD moments) with the same estimation code. The 2022–2023 estimation uses 2020-vintage ACS block groups; because the driving-duration matrix is built on 2010-vintage block groups, each 2020 block group is mapped to its nearest 2010-vintage neighbor by centroid distance (median drift 46m; 95th percentile 581m).

C.3.2 Consumer Demand Model

For the demand side, we do not have a bridge collapse in 2019 for identifying parameters, so we also compare observational estimates across periods. As above, we normalize by the constant term since parameter magnitudes in logit choice models are otherwise difficult to compare. The observational structural parameters are qualitatively similar, with $\hat{\beta}^d / \lvert\hat{\alpha}\rvert = 0.047$ in 2019 versus 0.107 in 2024.

D Supplementary Entry Model Analysis

This appendix presents supplementary analyses for the restaurant entry model, including isolation of the endogenous feedback channel and bootstrap confidence intervals.

D.1 Isolating the Endogenous Feedback Channel

We attribute the endogenous-market model's advantage to the feedback loop in which restaurant entry reshapes local household composition and thus future demand. To test whether this restaurant→population channel specifically drives the improvement, rather than other features of the household sorting model, we conduct two exercises that vary the household covariates driving population forecasts.

Table D1 — Isolating the Endogenous Feedback Channel

	Exogenous-Market		Endogenous-
	Dept. of Plan.	No Rest.	Market
HH Model Parameters			
$\hat{\beta}^{\text{rest}}$	—	not included	included
Out-of-Sample Fit (2018–2019)			
Total MSE	2.006	2.009	1.898
Improvement vs. Exogenous-Market			
Total MSE impr. over Dept. of Plan.	—	−0.16%	5.37%

Notes. Each endogenous-market column uses the household model to construct population forecasts, but varies which household parameters are included. The exogenous-market column uses Maryland Department of Planning population projections (Appendix D.5). "No Rest." constrains the household restaurant-proximity coefficients to zero while holding all other household coefficients (moving cost, distance, distance-to-CBD, same-county, non-restaurant business density, and log-density/log-income push–pull terms) at their headline estimates; "Endogenous-Market" uses the full model with restaurant proximity.

Results are reported in Table D1. Under the mechanism-null specification, in which the household restaurant-proximity coefficients are constrained to zero while all other household covariates are held fixed (column "No Rest."), total out-of-sample MSE rises to 2.009, essentially indistinguishable from the exogenous-market baseline of 2.006 and substantially worse than the full endogenous-market specification's 1.898. This collapses the endogenous-market wedge from 5.37% (full model) to −0.16% (mechanism-null), a decisive attribution of the improvement to the restaurant→population channel. The predictive gain disappears entirely when the restaurant→population feedback is removed, strongly suggesting that this feedback channel specifically drives the endogenous-market advantage rather than generic features of the household sorting model.

D.2 Household Utility Specification Robustness

To assess whether our baseline entry-model results are sensitive to the specification of the household utility function, we re-estimate the entry model under a maximally-sparse household specification that retains only the moving cost, restaurant proximity at 1 mile, restaurant proximity at 2 miles, and log driving-time decay — four parameters total, dropping all the auxiliary covariates (non-restaurant retail density, same-county fixed effect, log-density and log-income push/pull terms, and distance-to-CBD) that appear in the baseline.

Table D2 reports the household utility parameter estimates for both specifications, with 95% Bayesian bootstrap percentile confidence intervals in brackets for the baseline column (100 draws; see Appendix E.4). The sparse column reports point estimates only. The two specifications agree on the sign of the two key restaurant-proximity coefficients ($\hat{\beta}_1^{\text{rest}} < 0$, $\hat{\beta}_2^{\text{rest}} > 0$), consistent with the non-monotonic residential preference identified in Section 5.2.2. The sparse specification recovers a larger 1-mile disamenity coefficient than the baseline (-0.065 vs. -0.038) and a steeper distance decay (-5.5 vs. -4.2), reflecting that in the absence of auxiliary controls the distance and restaurant-proximity terms absorb residual variation that the baseline specification allocates to the push/pull, distance-to-CBD, and same-county covariates.

Downstream, when the entry model is estimated under the maximally-sparse household specification instead of the baseline, the total out-of-sample MSE improvement is 7.50%

Table D2 — Household Utility Parameter Comparison: Baseline vs. Maximally-Sparse Specification

	Baseline	Sparse
β^{move}	-6.708 [-6.711, -6.701]	-5.536
β_1^{rest}	-0.038 [-0.068, -0.009]	-0.065
β_2^{rest}	0.023 [0.018, 0.029]	0.027
$\beta^{\text{nonrest,0-1}}$	0.000 [-0.011, 0.008]	—
$\beta^{\text{nonrest,1-2}}$	-0.005 [-0.008, -0.001]	—
β^{dist}	-4.247 [-4.250, -4.239]	-5.495
$\beta^{\text{dist,cbd}}$	-0.064 [-0.065, -0.014]	—
$\beta^{\text{same-cty}}$	1.281 [1.276, 1.290]	—
$\beta^{\text{log-dens,push}}$	0.051 [0.029, 0.052]	—
$\beta^{\text{log-dens,pull}}$	0.149 [0.123, 0.150]	—
$\beta^{\text{log-inc,push}}$	-0.148 [-0.159, -0.147]	—
$\beta^{\text{log-inc,pull}}$	0.253 [0.242, 0.255]	—
N bootstrap draws	100	— (point estimate)

Notes. Household utility parameter estimates under the baseline (12 parameters, 100 bootstrap draws) and maximally-sparse (4 parameters, point estimate only) specifications. 95% Bayesian bootstrap percentile confidence intervals in brackets for the baseline column (see Appendix E.4). Dashes indicate parameters not included in the sparse specification. Both specifications use log driving-time as the distance metric.

(1.856 vs. 2.006), larger than the 5.37% under the baseline and highly statistically significant (Diebold–Mariano $z = 9.71$, $p < 0.001$). Our baseline choice to include the auxiliary controls is therefore a conservative one: without them, the endogenous-market advantage would be reported as larger, not smaller.[30]

D.3 Household Restaurant-Proximity Scaling Robustness

To directly probe how sensitive our downstream entry-model conclusions are to the magnitude of the estimated household restaurant-proximity coefficients, we re-estimate the entry model after multiplicatively scaling both $\hat{\beta}_1^{\text{rest}}$ and $\hat{\beta}_2^{\text{rest}}$ by 0.5 (severe attenuation) and 1.5 (severe inflation), holding all other parameters at their headline estimates. This bracket is much wider than plausible sampling uncertainty on either coefficient and thus stress-tests whether our

30. The counterfactual road-corridor rankings, in contrast, are more sensitive to the household specification; see the discussion in Section 6.3 for the baseline's rankings.

conclusions are driven by the specific magnitudes recovered from the zoning quasi-experiment.

Table D3 reports total out-of-sample MSE and the endogenous-market improvement at each scale. Both attenuation and inflation leave the endogenous-market advantage of comparable magnitude: the total-MSE improvement is 5.01% at 0.5× and 4.75% at 1.5×, compared with 5.37% at the point estimate. The direction and magnitude of the wedge are preserved across the ±50% bracket, confirming that our downstream conclusions are not an artifact of the specific magnitudes of the restaurant-proximity estimates.

Table D3 — Entry Model Robustness: Scaling Household Restaurant-Proximity Coefficients

Scaling factor	Total MSE	Improvement vs. exog	Endogenous $\hat{\mu}$
0.5× (severe attenuation)	1.905	5.01%	0.084
1.0× (point estimate)	1.898	5.37%	0.086
1.5× (severe inflation)	1.911	4.75%	0.034

Notes. Both $\hat{\beta}_1^{\text{rest}}$ and $\hat{\beta}_2^{\text{rest}}$ are multiplied by the scaling factor; all other household coefficients (moving cost, distance, distance-to-CBD, same-county, non-restaurant business density, and log-density/log-income push–pull terms) are held at their headline estimates. The exogenous-market total MSE (2.006) is unchanged across scales. Improvement is (2.006 − Total MSE)/2.006.

D.4 Counterfactual Details

This appendix provides additional details on the transit improvement counterfactual presented in Section 6.3.

D.4.1 Corridor Definitions and Simulation

The counterfactual simulates a 10% reduction in travel times for origin-destination pairs along a hypothetical road corridor. We adopt 10% as an illustrative magnitude rather than as a calibration to any specific real-world intervention; our exercise focuses on the relative comparison between the endogenous-market and exogenous-market specifications, rather than on the absolute magnitudes of the entry responses. Each corridor is defined by an origin coordinate, a destination coordinate, and a buffer radius ranging from 1.5 to 2.5

miles depending on corridor length. An OD pair is affected if both the origin CBG and the destination restaurant location fall within the buffer distance of the corridor centerline. Travel times for affected OD pairs are reduced by 10%, while all other OD pairs retain their baseline travel times.

The travel-time reduction enters through the demand model of Section 4, since shorter travel times increase the convenience of near-road restaurants, shifting consumer spending toward them and raising local market size M_{jt}. We then re-solve for equilibrium CCPs under the modified demand.[31]

We simulate corridors with the origin and destination coordinates listed in Table D4's corridor labels; exact centerline endpoints are documented in the replication package. The three corridors highlighted in Section 6.3.1 (Urban Link, Downtown Connector, and Suburban Link) span three levels of existing local catchment density: moderate-density with growth headroom (Urban Link), high-density inner Baltimore (Downtown Connector), and sparse-catchment suburban (Suburban Link).

D.4.2 Welfare Decomposition by Corridor

Table D4 reports the change in restaurant entry (ΔEntry, i.e., the percentage change in gross new entrants over the simulation horizon) and in total firm surplus (ΔFS) for each of the three main-text corridors, under both the endogenous-market and exogenous-market specifications. We report ΔFS as a supporting welfare check on the entry story of Section 6.3; note that ΔFS and ΔEntry capture distinct margins and need not co-move, since firm surplus aggregates over all firms (existing and new) while entry counts only new firms, so a densely-served corridor can register a positive ΔFS as remaining firms absorb the added visits even when ΔEntry is flat or slightly negative.

Firm surplus tracks the entry-side story in direction across all three corridors. Under

31. We iterate the counterfactual CCPs to a cap of 20, at which the baseline's predicted base-state (508 entries, 13,740 exits over the simulation horizon) matches the original NPL-estimated base-state. To ensure a fair comparison, we also iterate the baseline CCPs the same number of times under unmodified travel times, and compute entry changes as the difference between the counterfactual and re-iterated baseline.

Table D4 — Counterfactual Road Improvement: Entry and Firm Surplus by Corridor

	Endogenous-Market		Exogenous-Market	
Corridor	ΔEntry%	ΔFS	ΔEntry%	ΔFS
Urban Link	0.247	16.1	0.189	12.8
Downtown Connector	0.149	29.2	0.040	3.0
Suburban Link	0.058	21.7	0.255	14.6

Notes. Counterfactual effects of a 10% travel-time reduction along each corridor. ΔEntry is the percentage change in gross metro-area restaurant entry (new entrants over the simulation horizon). ΔFS is the change in total firm surplus, aggregated over all firms (remaining and new) and expressed in the model's private-value units (per-period integrated V scaled by the profit shock standard deviation σ). ΔFS is not in dollars and should not be interpreted as a social-welfare measure; it summarizes the private-value component of the entry model's payoff and is intended here as a directional cross-check on the entry-side story, not as a policy-ranking metric on its own. The three main-text corridors are also shown in Figure 5.

the "Urban Link" both specifications project firm-surplus gains, with the endogenous specification larger (16.1 vs. 12.8) in line with the endogenous-market entry premium. Under the "Downtown Connector", the endogenous specification records the largest firm-surplus gain of the three corridors (29.2 vs. 3.0 under the exogenous specification, a nearly tenfold understatement)—the strong local sorting response inside a denser corridor passes through to incumbent per-firm profitability that a naive projection cannot see. Under "Suburban Link" both specifications also project firm-surplus gains, and here the endogenous specification is *larger* on firm surplus (21.7) than the exogenous specification (14.6) despite recording only a small entry response (0.058%): because the local catchment is sparse, few new entrants are drawn in and the transit-induced demand boost is captured by the existing incumbents through per-firm profit rather than through new entry.

D.5 Department of Planning Population Projection Data

The exogenous-dynamics baseline uses population projections from the Maryland Department of Planning (MDP). MDP publishes population projections at five-year intervals for each of Maryland's 24 jurisdictions (23 counties plus Baltimore City). Our study area spans 10 counties: Anne Arundel, Baltimore City, Baltimore County, Carroll, Harford, Howard, Kent, Montgomery, Prince George's, and Queen Anne's.

We use MDP projection anchor points (2000, 2010, 2020, 2025, 2030) and linearly interpolate population levels between adjacent anchors. For each simulated period t and each county c, the population growth factor g_{ct} is computed as the year-over-year ratio of these interpolated populations,

$$g_{ct} = \frac{P_{c,t+1}}{P_{c,t}} - 1, \tag{D.1}$$

where $P_{c,t}$ is the linearly-interpolated MDP population for county c at year t. In each simulated period we multiply every CBG's population share by its county factor $(1 + g_{ct})$ and renormalize shares to sum to one across the metro area; market size M_{jt} is then recomputed from the updated population through the demand model. This imposes the official county-level aggregate trend while preserving within-county heterogeneity in market size.[32]

32. For further documentation, see Maryland Department of Planning, "Projections" (https://planning.maryland.gov/MSDC/Pages/projection/projections.aspx).

E Model Details

This appendix provides technical details on estimation procedures for the three models developed in the main text.

E.1 Consumer Demand Model

E.1.1 Model Primitives

Consumers choose among J restaurants plus an outside option. The deterministic component of utility for a consumer in CBG g visiting restaurant j is:

$$V_{gj} = \alpha + \beta^d \cdot \log(d_{gj}) \tag{E.1}$$

where α is a constant and d_{gj} is driving time in minutes from CBG g to restaurant j. Total utility is $V_{gj} + \varepsilon_{gj}$, where ε_{gj} is an i.i.d. Type I extreme value shock. The outside option deterministic utility is normalized to zero: $V_{g0} = 0$.

Choice probabilities follow the logit formula:

$$P_{gj} = \frac{\exp(V_{gj})}{1 + \sum_k \exp(V_{gk})} \tag{E.2}$$

E.1.2 Aggregation from Visits to Spending

The model predicts "visitors" to each restaurant:

$$\text{visitors}_j = \sum_g \text{Pop}_g \cdot P_{gj} \tag{E.3}$$

We convert model-predicted visitors to spending using mean spend per transaction for each restaurant:

$$\bar{m}_j = \frac{\sum_t \text{spend}_{jt}}{\sum_t \text{transactions}_{jt}} \tag{E.4}$$

Model-predicted spending is then $\hat{\text{spend}}_j = \text{visitors}_j \cdot \bar{m}_j$. The DiD estimand compares changes in predicted spending between treated and control restaurants.

E.1.3 Estimation

We estimate (α, β^d) using two moment conditions, with α concentrated out via an inner loop. Pre-collapse travel times d_{gj}^{pre} and post-collapse travel times d_{gj}^{post} (augmented with congestion multipliers from Chang et al. 2025) are computed once from the OSRM network. The estimation then proceeds as an outer search over β^d:

1. For each candidate β^d, solve for $\alpha(\beta^d)$ via an inner one-dimensional optimization that sets mean pre-period predicted transactions equal to mean observed transactions: $g_1(\alpha, \beta^d) = \frac{1}{J}\sum_j [\hat{y}_j^{\text{pre}}(\alpha, \beta^d) - y_j^{\text{pre}}] = 0$
2. Given $\alpha(\beta^d)$, compute choice probabilities P_{gj}^t for $t \in \{\text{pre}, \text{post}\}$ and aggregate to model-predicted spending $\hat{y}_j^t \cdot \bar{m}_j$
3. Run the binary-treatment DiD regression on model-predicted spending to obtain $\widehat{\text{DiD}}^{\text{model}}(\alpha(\beta^d), \beta^d)$
4. Choose β^d to match the model-implied DiD to the empirical estimate

This concentration approach yields a unique $\alpha(\beta^d)$ for each β^d, reducing the problem to a one-dimensional root-finding exercise. The outer loop identifies β^d by matching the model-implied DiD coefficient to the empirical estimate:

$$g_2(\beta^d) = \widehat{\text{DiD}}^{\text{model}}(\alpha(\beta^d), \beta^d) - \widehat{\text{DiD}}^{\text{data}} = 0 \tag{E.5}$$

where $\widehat{\text{DiD}}^{\text{model}}$ is computed by running the same binary-treatment regression on model-predicted spending $\hat{y}_j \cdot \bar{m}_j$, with $\bar{m}_j$ the observed mean spend per transaction. This concentration approach is equivalent to jointly minimizing $Q(\alpha, \beta^d) = g_1^2 + g_2^2$ with the identity weighting matrix, since the model is exactly identified.

The difference-in-differences exploits the March 26, 2024 collapse of the Francis Scott Key Bridge. We define treatment based on driving duration to the bridge location computed via OSRM:

- **Treated restaurants**: Within 30 minutes driving time of the bridge
- **Control restaurants**: Beyond 30 minutes driving time (including the 30–45 minute intermediate zone as controls)
- **Pre-period**: July 2023–March 2024 (9 months)
- **Post-period**: April–December 2024 (9 months)

We use a symmetric 9-month window around the collapse, matching the reduced-form DiD specification in column (3) of Table 1, with monthly congested travel time matrices for each post-collapse month. Including the intermediate zone as controls maximizes statistical power; the reduced-form DiD estimates are nearly identical whether this intermediate zone is included or excluded (Table 1).

Standard errors are computed via Bayesian bootstrap with 100 replications; see Section E.4 for details.

E.1.4 Alternative Estimation via Orthogonality

For comparison, we also estimate the model by replacing the DiD moment with a cross-sectional orthogonality condition that requires spending prediction errors to be uncorrelated with log travel duration. Both moments are computed using only pre-collapse months ($t \in \mathcal{T}_{\text{pre}}$):

$$g_1(\alpha, \beta^d) = \frac{1}{|\mathcal{T}_{\text{pre}}| \cdot J} \sum_{t \in \mathcal{T}_{\text{pre}}} \sum_{j=1}^{J} r_{jt} = 0 \tag{E.6}$$

$$g_2(\alpha, \beta^d) = \frac{1}{|\mathcal{T}_{\text{pre}}| \cdot J} \sum_{t \in \mathcal{T}_{\text{pre}}} \sum_{j=1}^{J} r_{jt} \cdot \log d_j = 0 \tag{E.7}$$

where $r_{jt} = y_{jt} - \hat{y}_{jt}(\alpha, \beta^d)$ is the spending prediction residual at restaurant j in month t, and d_j is the population-weighted average pre-period travel duration to restaurant j. The first moment pins down α by requiring zero average residual; the second identifies β^d by requiring that residuals be orthogonal to log travel duration in the cross-section. Unlike the DiD-based moment, this condition does not exploit variation from the bridge collapse. This estimator yields $\hat{\beta}^d = 2.08$, which differs severely from the DiD-identified estimate (-0.39); we discuss the interpretation in Appendix C.1.

E.2 Household Location Choice Model

Given candidate parameters $\theta = (\beta^{\text{move}}, \beta_1^{\text{rest}}, \beta_2^{\text{rest}}, \beta^{\text{dist}})$, we simulate population dynamics across 3,106 CBGs over 7 years (2013–2019). A household currently in CBG g obtains utility $U_{ggt} = V_{gt} + \varepsilon_{ggt}$ from staying and $U_{gg't} = \beta^{\text{move}} + V_{g't} + \beta^{\text{dist}} \cdot d_{gg'} + \varepsilon_{gg't}$ from moving to $g' \neq g$, where V_{gt} is the flow utility defined in equation (8). Because amenities enter both the stay and move utilities, locations with high V_{gt} both attract in-migrants and retain current residents. The simulation proceeds as follows:

1. Initialize population at 2013 levels from Census data.[33]
2. For each year t, compute location utilities V_{gt} using the year-t cumulative counts of restaurants within 1 mile and within 2 miles of each CBG centroid, built annually from the Data Axle entry/exit panel.
3. Compute choice probabilities $P_{gg',t}$ from the logit formula (including $P_{gg,t}$, the stay probability).
4. Update population: $\text{Pop}_{g,t+1} = \sum_{g'} \text{Pop}_{g't} \cdot P_{g'g,t} + \text{ExtMig}_{gt}$.
5. Repeat for $t = 2013, \ldots, 2018$.

The external migration term $\text{ExtMig}_{gt} = \text{Pop}_{gt} \cdot r_{c(g),\text{period}}$ applies a county-specific net migration rate $r_{c(g),\text{period}}$ proportionally to each CBG's current population, which is equivalent to distributing the county-level external migration total across CBGs in proportion to their share of county population. The rate $r_{c(g),\text{period}}$ is the per-capita annual net inflow into county $c(g)$ from areas outside the 10-county study area, calibrated from Census ACS migration data

33. The simulation initializes at the observed 2013 population distribution and evolves population shares forward via logit choice probabilities; we do not separately recover or rationalize unobserved CBG quality ξ_g. Because our moments target *changes* in population (via DiD coefficients), identification works on how marginal restaurant-density shifts cause population reallocation conditional on the 2013 baseline. Simulated counterfactual populations should accordingly be interpreted as marginal deviations from the 2013 initial condition rather than as standalone level forecasts.

with separate rates for the pre-treatment period (2011–2015 ACS vintage) and post-treatment period (2015–2019 ACS vintage); rates range from approximately -1.8% to 3.0% per year across counties and periods.

We target four types of moments:

1. **DiD moments**: Running the reduced-form regression on simulated data and matching coefficients to empirical targets ($\hat{\beta}^{\text{data,0-1}} = -5.14$ for the 0–1 mile ring, $\hat{\beta}^{\text{data,1-2}} = 10.44$ for the 1–2 mile ring; from Table 3). These identify β_1^{rest} and β_2^{rest}.
2. **Mobility rate**: Annual share of households who move (target: 9.68% from Census ACS). Identifies β^{move}.
3. **Same-county rate**: Share of movers who relocate within the same county (target: 76.1% from Census ACS). Identifies β^{dist}.
4. **County flow shares**: The county-to-county flow shares s_{od}^{data}, where s_{od}^{data} is the fraction of within-CBSA movers leaving origin county o who relocate to destination county d (the 84 origin–destination pairs with non-zero observed migration, out of 10×10 possible county pairs including same-county stays where $o = d$; pairs with no observed ACS flow are dropped). Empirical targets are computed from Census ACS county-to-county migration tables. These serve as overidentifying restrictions, providing additional discipline on the model's spatial migration patterns.

The GMM objective function is:

$$Q(\theta) = \sum_m w_m \left(g_m^{\text{sim}}(\theta) - g_m^{\text{data}} \right)^2 \tag{E.8}$$

where w_m are moment weights, set heuristically as the smallest values that ensure near-exact matches between simulated and sample moments in practice. Specifically, we set $w = 100$ for the DiD moments and $w = 10{,}000$ for the mobility rate and same-county rate, leading to approximately equal effective weight across the four primary moments, and $w = 5$ per county-pair flow moment, again set heuristically as a positive weight that does not induce any salient divergence in the primary moments. The distance preference parameter β^{dist} is measured per 10 miles.

E.3 Restaurant Entry Model

This section provides complete computational details for the dynamic entry model estimated via nested pseudo-likelihood (NPL) in Section 6.

E.3.1 Profit Function

Flow profits for an incumbent at location j in period t are:

$$\pi_{jt} = \mu \cdot \frac{M_{jt}}{n_{jt}} \tag{E.9}$$

where μ scales per-restaurant demand revenue under equal sharing, M_{jt} is market size (consumer demand aggregated from CBG origins, rescaled by 1/1,000 as a numerical normalization that is absorbed into the units of μ), and n_{jt} is competitor count. When $n_{jt} = 0$, we set $M_{jt}/n_{jt} = 0$.

E.3.2 State Space and Equilibrium Concept

Following the aggregate-state Markov perfect equilibrium literature, including the oblivious equilibrium of Weintraub, Benkard, and Van Roy (2008) and the moment-based Markov equilibrium of Ifrach and Weintraub (2017), we specialize to a fully symmetric framework in which each firm conditions on the low-dimensional aggregate state (n_{jt}, M_{jt}) rather than on the full vector of rival identities and characteristics. The aggregate transition $F(x_{t+1} \mid a_t, x_t)$ treats all firms in n_{jt} symmetrically, with each incumbent exiting with probability $p^{\text{exit}}(x_{jt})$ and each potential entrant entering with $p^{\text{enter}}(x_{jt})$, both functions of the aggregate state.

This has distinct implications for the two value functions. For an incumbent's continue/exit decision, symmetry implies that the focal firm's individual choice enters as one of n_{jt} exchangeable contributions to the aggregate transition, since the aggregate Binomial transition treats the focal and other incumbents identically. The forward simulation described below accordingly draws the focal firm's own exit independently from the same p^{exit} that governs the aggregate. For an entrant's decision, the 1 contribution to $n_{j,t+1}$ enters the value function directly, as the entrant is by construction not in n_{jt}; the polynomial sieve described below evaluates the entrant's continuation value at the post-entry state $(n_{jt}+1, \log M_{j,t+1})$, with $M_{j,t+1}$ reflecting CCP-implied background industry evolution alongside the household response to the marginal entrant.

We note that this aggregate-state symmetric framework is exact in the limit of large industry size and a controlled approximation otherwise; identification draws on the cross-sectional pool of 3,106 CBGs and the within-location time variation across 2013–2019.

E.3.3 Initial CCP Estimation

We estimate flexible initial conditional choice probabilities (CCPs) from the data using polynomial logistic regression; these serve as the starting beliefs about future industry evolution that the NPL iteration below updates iteratively toward the model's fixed point. The state for each location-period is $(n_{jt}, \log M_{jt})$. Exit probabilities are estimated via binomial logistic regression on the subsample of locations with $n_{jt} > 0$, using degree-3 polynomial features with interactions. Entry probabilities are estimated via Poisson-derived logistic regression on all locations. Both estimators use L2 regularization ($\lambda = 10^{-4}$, excluding the intercept) for numerical stability.

E.3.4 Value Function Simulation

The integrated incumbent value $V(x_{jt})$ defined in equation (16) admits the forward-sum representation

$$V(x_{jt}) = \mathbb{E}_t\Bigg[\sum_{\tau=0}^{T} \delta^{\tau}\Big(\pi_{j,t+\tau} + \sigma\,\varphi\Big(\Phi^{-1}(1-p^{\text{exit}}_{j,t+\tau})\Big)\Big)\mathbf{1}\{\text{alive at } \tau\} + \frac{\delta^{T+1}}{1-\delta}\pi_{j,t+T}\mathbf{1}\{\text{alive at } T\} + \delta^{\tau_{\text{exit}}+1} SV\Bigg], \tag{E.10}$$

where $T = 20$ is the simulation horizon. The first term in the inner sum is the discounted flow profit while alive; the second is the per-period Aguirregabiria and Mira (2007) selected-shock contribution $\sigma\,\varphi(\Phi^{-1}(P_{\text{stay}}))$ accumulated at the same period discount as the profit; the third is the terminal perpetuity extending profits beyond T; and τ_{exit} is the period in which the firm draws its exit, receiving scrap value SV at the next-period discount.

Because all components are linear in (μ, SV, σ), V in equation (E.10) is linear in each. We exploit this by decomposing V into pre-computable basis components. Specifically, for each location j and period t, we simulate 200 Monte Carlo paths of industry evolution over a $T = 20$ horizon. In each path, we sample exits and entries per period using the current CCPs, evolve population via the household model (or Department of Planning projections for the exogenous specification), and update market size M.[34] Along each path, the focal firm's own exit is also drawn stochastically from the same CCPs governing aggregate exit; profits accumulate only while the focal firm is alive, and upon exit the firm receives the scrap value SV at the next-period discount factor. We accumulate the discounted profit basis component along each path:

$$\begin{aligned} V_{jt}^{M/n} = &\sum_{\tau=0}^{T} \delta^{\tau} \frac{M_{j,t+\tau}}{n_{j,t+\tau}} \cdot \mathbf{1}_{\text{alive at } \tau} \\ &+ \frac{\delta^{T+1}}{1-\delta} \frac{M_{j,t+T}}{n_{j,t+T}} \cdot \mathbf{1}_{\text{alive at } T} \end{aligned} \tag{E.11}$$

The terminal value assumes the period-T state persists indefinitely if the focal firm is still active at T. We average over draws to obtain $\bar{V}_{jt}^{M/n}$, $\bar{V}_{jt}^{\text{exit}}$, and $\bar{V}_{jt}^{\varphi}$, so that the integrated incumbent value is $V(x_{jt}) = \mu \cdot \bar{V}_{jt}^{M/n} + SV \cdot \bar{V}_{jt}^{\text{exit}} + \sigma \cdot \bar{V}_{jt}^{\varphi}$ for any choice of (μ, SV, σ) (with σ normalized to 1).

34. The forward projection omits external migration, which enters the household model estimation but would require restaurants to forecast future county-level migration from outside the study area. Net external migration accounts for roughly 3–6% of total within-study-area population flows across the two ACS periods we use (2011–2015 and 2015–2019).

E.3.5 Post-Action Value Approximation

To compute the action-specific value differences that drive the probit thresholds in equation (19), we need the expected next-period integrated value at the post-action state for both stay and enter actions: $v^{\text{cont}}(x_{jt}) = \mathbb{E}[V(x_{j,t+1}) \mid \text{stay}, x_{jt}]$ and $V^{\text{ent}}(x_{jt}) = \mathbb{E}[V(x_{j,t+1}) \mid \text{enter}, x_{jt}]$. We approximate both using a polynomial sieve fitted to the forward-simulated V basis components.

For each period t, we fit degree-3 polynomials of $(n, \log M)$ to each basis component (the profit basis $\bar{V}_{jt}^{M/n}$, the scrap-weight basis $\bar{V}_{jt}^{\text{exit}}$, and the Aguirregabiria and Mira (2007) selected-shock basis $\bar{V}_{jt}^{\varphi}$), and then evaluate these polynomials at two post-action states. For the post-entry state we use $(n+1, \log M'_{\text{enter}})$, where M'_{enter} is the demand-model-corrected market size accounting for the household response to the marginal entrant. For the post-stay state we use $(n, \log M'_{\text{stay}})$, where M'_{stay} reflects one-step household evolution under unchanged restaurant count. The household corrections use an "expected next state" construction, applying one CCP-predicted industry step before evaluating V at the post-action state. Because each basis component is linear in (μ, SV, σ), the resulting v^{cont} and V^{ent} remain linear in these parameters, preserving the tractability of the inner pseudo-likelihood optimization.

The background industry state used inside the focal entrant's continuation value is held at the observed state: we add the focal entrant to the current restaurant stock without advancing the background stock, since the forward-simulated value already prices in expected industry transitions and adding another partial expected transition at the entry step would double-count.

E.3.6 Entry and Exit Probabilities

Entry and exit decisions follow probit threshold rules with firm-specific shocks. Substituting the action-specific values from equation (17) into the choice differentials (recalling that current-period π_{jt} cancels in $v^{\text{stay}} - v^{\text{exit}}$), the exit probability for each incumbent is

$$p_{jt}^{\text{exit}} = 1 - \Phi\left(\frac{\delta\left(v^{\text{cont}}(x_{jt}) - SV\right)}{\sigma}\right), \tag{E.12}$$

and the entry probability for each potential entrant is

$$p_{jt}^{\text{enter}} = \Phi\left(\frac{\delta\, V^{\text{ent}}(x_{jt}) - EC}{\sigma}\right), \tag{E.13}$$

where σ is the shock standard deviation (normalized to 1) and EC is paid at the time of the entry decision. Exit counts follow a binomial distribution with n_{jt} trials and probability p_{jt}^{exit}. Entry counts follow a Poisson distribution with rate $\lambda_{jt} = K \cdot p_{jt}^{\text{enter}}$, where $K = 100$ is the pool of potential entrants, set well above the maximum observed entry counts per location-year so that the bound is non-binding.[35]

E.3.7 Identification

Identification of $\theta = (\mu, EC, SV)$ follows familiar lines for entry-exit models (Aguirregabiria and Mira 2007; Aguirregabiria and Suzuki 2014); see Section 6.1.1 for the main-text discussion of how variation in entry and exit rates separately identifies each parameter. Cross-sectional variation across 3,106 CBGs and within-location time variation across 2013–2019 jointly identify the parameter combination, though as noted by Aguirregabiria and Suzuki (2014) and reflected in the bootstrap distribution reported in Section 6.1.2, individual coordinates trade off along a manifold and are jointly rather than separately identified at the marginal level.

E.3.8 Structural Estimation

We normalize $\sigma = 1$ without loss of generality, since only the ratios V/σ, EC/σ, and SV/σ affect choice probabilities. The three estimated parameters are $\theta = (\mu, EC, SV)$, interpreted

35. The Poisson approximates the binomial when K is large and p^{enter} is small, which holds in our setting where most locations see zero or one entry per year.

as measured in units of σ. We maximize the pseudo log-likelihood using L-BFGS-B with 25 multi-start initializations. μ and EC are constrained to be positive; SV is unconstrained. Confidence intervals are computed via Bayesian bootstrap (Section E.4).

The full estimation iterates the following NPL loop: (1) estimate initial CCPs from the data; (2) simulate value function basis components forward under these CCPs; (3) maximize the pseudo log-likelihood to recover $\hat{\theta}$; (4) compute the structural model's implied CCPs using $\hat{\theta}$ and the simulated values; (5) update the CCPs to $\text{CCP}^{(k+1)} = \text{CCP}^{(k)}_{\text{structural}}$ and repeat from step (2) until the CCPs are consistent with the structural model's implied probabilities.

E.4 Bootstrap Inference

We compute confidence intervals using the Bayesian bootstrap (Rubin 1981), which has two advantages over the classical nonparametric bootstrap in our setting. First, it avoids the discreteness of resampling with replacement, producing smoother sampling distributions with a coherent posterior interpretation under a nonparametric prior (Newton and Raftery 1994). Second, and more importantly for discrete choice models, it preserves the choice set structure: rather than resampling choice occasions (which would alter the set of alternatives available), we reweight contributions to the likelihood while keeping the choice sets intact.

For each bootstrap replication, we draw $p \sim \text{Dirichlet}(\alpha, \ldots, \alpha)$ over the n resampling units with concentration parameter $\alpha = 4$ (Shao and Tu 1995) and form mean-one weights $w_i = n\, p_i$ to reweight the objective before re-estimating the model. These weights have variance $(n-1)/(n\alpha+1) \approx 1/\alpha$ for large n, smaller than the unit variance of the multinomial weights implicit in the classical nonparametric bootstrap, so we rescale by the variance-restoring factor $c = \sqrt{(n\alpha+1)/(n-1)} \approx \sqrt{\alpha}$. For the demand and household models we report the rescaled bootstrap standard error (parenthetical values), equal to c times the raw bootstrap standard deviation. For the entry model, where propagation through the nonlinear value-function simulation can produce asymmetric sampling distributions, we report rescaled 95% percentile intervals (bracketed values), inflating the raw percentile interval about its median so that the rescaled q-th percentile is $\text{med} + c\,(\hat{\theta}^*_{(q)} - \text{med})$, with $\hat{\theta}^*_{(q)}$ the raw q-th bootstrap percentile and med the bootstrap median (cf. Shao and Tu 1995).

For the consumer demand model, we draw Dirichlet weights over the J restaurants and weight the moment conditions in the GMM objective ($B = 100$ replications). For the household location choice model, we draw weights over CBGs and weight the moment conditions ($B = 100$ replications). In both cases, each replication recomputes the DiD targets from the reweighted data before re-estimating the structural model, so the resulting confidence intervals fully account for sampling uncertainty in the quasi-experimental first stage. For the restaurant entry model, each of $B = 100$ replicates draws Dirichlet weights ($\alpha = 4$) over CBGs, redraws an upstream household and demand parameter vector from the saved upstream bootstrap chains, and re-estimates the full nested pseudo-likelihood, propagating estimation uncertainty from both upstream stages through to the entry-model structural parameters. The joint Wald test on $\theta = (\mu, EC_0, SV_0)$ against $H_0 : \theta = 0$ yields $W = 67{,}335$ ($p < 0.001$ with χ^2_3).