CouncilHound Development Impact Methodology

Hunter Brown
huntermbrown@gmail.com
CouncilHound — impact analysis subsystem

August 2026
model version: joint destination–mode Huff, census-block weights;
cost defaults calibrated to city staff estimates

Abstract

CouncilHound produces screening-level impact estimates for proposed local development projects. A deterministic economic module (spatial-interaction retail capture, walking activity, an employment ledger) and a deterministic fiscal module (revenue and service-cost ranges) generate every number; a language model is used only to extract project facts from public documents under a verbatim-quote firewall and to write a narrative that a validator constrains to the structured outputs. This note documents each displayed metric's formula, its academic or statistical basis, the assumptions that drive its uncertainty range, and the implementation decisions behind the current system. The estimates rank likely orders of magnitude under stated assumptions; they are decision-support evidence, not forecasts.

1 Scope and design principles

1. **Deterministic quantities.** All numbers are produced by closed-form models over open data. The language model never supplies a quantity: extracted project facts require a verbatim source quote ($\leq$15 words) that is verified against the document text in code, and fields without one are nulled. The narrative report is post-validated so that every number it contains matches a published metric, a stated assumption, a provenance string, or a deterministic module note.
2. **Assumption transparency.** Every contestable parameter is a named assumption with a central value, low/high bounds, and a recorded basis. Bounds propagate through all formulas by interval arithmetic, so each metric carries a range derived from the same assumptions it cites.
3. **Fail loudly.** Missing inputs (an unpinned tax rate, an absent assessment record) degrade the affected metric to "not computed — reason," never to a silent default.
4. **Screening honesty.** Capture, walking, and foot-traffic figures are screening estimates in the spatial-interaction tradition; the report states this in plain language and names the most sensitivity-driving assumptions.

2 Notation and data inputs

U	proposed dwelling units (extracted from project documents)
ρ	occupancy rate at stabilization (`occupancy_rate`)
$\bar{s}$	persons per occupied multifamily unit (`avg_hh_size_multifamily`)
y	mean household income of the site's census tract (ACS B19025/B11001)
π	new-construction income premium (`income_premium_new_construction`)
e_c	CES average annual expenditure for category c (BLS CES 2023)
y_{CES}	CES average pre-tax income per consumer unit ($101,805, 2023)
ε_c	expenditure–income elasticity for category c (Engel scaling)
κ	CES scaling factor (`ces_scale`)
A_j	attractiveness of destination j (POI count = 1 per business)
$t_j^w,\ t_j^d$	walk / drive travel time (minutes) from the site to j
$\beta_w,\ \beta_d$	walk / drive impedance decay per minute
w_c	zero-impedance walk preference for category c
τ	daily walk trips per resident (`walk_trips_per_resident_day`)

Abbreviations used throughout: ACS is the Census Bureau's American Community Survey; CES is the U.S. Bureau of Labor Statistics (BLS) Consumer Expenditure Survey [13]; NHTS is the Federal Highway Administration's (FHWA) National Household Travel Survey [14]; POI is point of interest.

Networks are OpenStreetMap (OSM) walk and drive graphs built with OSMnx [9]; walk edge times use 4.8 km/h. Population weights use 2020 Census Public Law (PL) 94-171 block centroids. Points of interest merge Overture Maps, OpenStreetMap, and (when licensed) Foursquare Open Source Places into a nine-class taxonomy. Travel times come from single-source Dijkstra runs over edge minutes; each business inherits the time of its nearest network node.

3 Demand module

3.1 New households and residents

$$N_{hh} = U \cdot \rho \tag{1}$$
$$N_{res} = N_{hh} \cdot \bar{s} \tag{2}$$

Occupancy at stabilization is benchmarked against the Census Housing Vacancies and Homeownership (HVS) series [18]. Household size uses a multifamily (MF) bedroom-mix demographic multiplier in the tradition of Listokin et al. [12] ($\bar{s} = 1.8$, range 1.5–2.2, for a typical studio/one-bedroom/two-bedroom mix): the ACS B25010 tenure-level renter average (≈ 3.1 locally) was evaluated and rejected as an estimator because it is dominated by single-family renter households and is blind to unit mix. When a project's bedroom mix is known, the multiplier should be set from it (§10).

3.2 Aggregate income and category spending

$$Y = N_{hh} \cdot y \cdot \pi, \qquad r = \frac{y\,\pi}{y_{\text{CES}}} \tag{3}$$
$$S_c = N_{hh} \cdot e_c \cdot r^{\,\varepsilon_c} \cdot \kappa \tag{4}$$

The tract income y is the *mean* (ACS aggregate household income B19025 divided by households B11001, summed over the site tract's block groups) — the correct estimator for total purchasing power, since median $\times$ count understates aggregate income under right-skewed distributions. When the aggregate tables are suppressed the model falls back to the population-weighted median and records the substitution. Category spending uses BLS Consumer Expenditure Survey 2023 levels [13] scaled by the income ratio r raised to a per-category expenditure elasticity ε_c (grocery 0.45, restaurant/bar 0.85, comparison 0.95, convenience 0.60, personal services 0.90, entertainment 1.05), consistent with the Engel gradients in the CES income-quintile tables: necessity spending grows sublinearly with income, so a high-income tract no longer implies proportionally more grocery spending.

4 Retail capture: joint destination–mode choice

Retail capture follows the probabilistic trade-area tradition of Huff [1, 2] with the exponential impedance of the entropy-maximizing spatial-interaction family [3]. Destination and mode are chosen *jointly*, as a multinomial logit [4, 5] over (business, mode) pairs. For category c, business j, and mode $m \in \{w, d\}$:

$$P_c(j,m) \;=\; \frac{A_j^{\alpha}\,\omega_{c,m}\,e^{-\beta_m t_j^m}}{\sum_k A_k^{\alpha}\left(\omega_{c,w}\,e^{-\beta_w t_k^w} + \omega_{c,d}\,e^{-\beta_d t_k^d}\right)}, \qquad \omega_{c,w} = w_c, \quad \omega_{c,d} = 1 - w_c, \tag{5}$$

with $\alpha = 1$. The walk impedance $\beta_w = 0.10$ min^{-1} is anchored to empirical negative-exponential decay fits: purpose-specific walking estimates of 0.09–0.11 min^{-1} (shopping 0.094, restaurant 0.093) from Twin Cities travel-survey data [15], and 0.073 min^{-1} from the national 2009 NHTS [16]; the sensitivity band (0.05–0.15) spans the national-sample floor to a hostile-pedestrian-environment ceiling. The mode preference w_c has a clean interpretation: at zero travel time for both modes, the walk share is exactly w_c; as walk time grows, walking loses utility to driving and the trip (and its dollars) arrives by car instead. Annual capture and walk-arriving capture at business j:

$$C_j = \sum_c S_c \sum_m P_c(j,m), \qquad C_j^{walk} = \sum_c S_c\, P_c(j,w). \tag{6}$$

Two invariants are pinned by unit tests: $\sum_j C_j = \sum_c S_c$ (joint probabilities sum to one), and $\sum_j C_j^{walk} \le \sum_c S_c w_c$ with equality only when every destination is at zero travel time — the walk total is *endogenous*, not an assumption restated.

Own-retail destination. The project's ground floor competes inside the model rather than being appended outside it. Its retail square footage is converted to POI-count-equivalent attractiveness (A_{own} = sqft/`own_retail_sqft_per_equiv_poi`), distributed across a fixed category mix (50% restaurant/bar, 15% comparison, 15% convenience, 20% personal services), at zero travel time from the site. Huff's original attractiveness variable was selling area [1]; competitor attractiveness here is POI count because open POI data lacks floor area, so the own-retail conversion reconciles the two units.

In-city shares. For food-away and for all-retail, the in-city capture share is the probability mass at destinations inside the city boundary, $\theta = \sum_{j\in\text{city}} \sum_m P(j,m)$. These shares are the cross-module inputs to the meals-tax and sales-tax lines.

Reporting rollup. Businesses are grouped for display by DBSCAN (density-based spatial clustering of applications with noise) [8] ($\varepsilon = 150\,\mathrm{m}$, minPts $= 3$) and each cluster is labeled by the named street intersection nearest its centroid (from the OSM walk graph, within $\sim$450 m), falling back to the central member business where no named intersection is close; cluster values are sums of member-business captures and play no role in the allocation itself.

5 Walking activity

5.1 Resident walk trips and exact marginal flows

Daily walk trips are $T = N_{res} \cdot \tau$. These all-purpose trips are allocated over POI-weighted walk-network nodes with exponential decay and routed along shortest paths — deterministically, with no sampling — giving per-street-segment flows:

$$f_e \;=\; \sum_v T \cdot \frac{g_v\, e^{-\beta_w t_v}}{\sum_u g_u\, e^{-\beta_w t_u}} \cdot \mathbb{1}[e \in \mathrm{SP}(\text{site} \to v)]\,, \tag{7}$$

where g_v is the POI weight at node v and SP is the shortest path. The change caused by the project *is* the trips its residents generate, so no baseline differencing (and no sampling noise) enters the numerator.

5.2 Baseline index and percent comparison

The comparison baseline is a seeded, sampled, weighted betweenness index [6, 7] of today's walkers: $k = 2000$ origin–destination pairs drawn with origins proportional to census-block population within 400 m of each node and destinations proportional to POI weight, routed by shortest paths. Scaling the sampled counts by (total population weight $\times\ \tau)/k$ expresses the baseline in the same trips units, and the trip rate τ cancels from the percentage:

$$\Delta_e \;=\; 100 \cdot \frac{f_e}{n_e \cdot \bar{P}\,\tau/k} \quad \text{(reported only where } n_e \geq 5\text{)}, \tag{8}$$

with n_e the sampled baseline count on edge e and $\bar{P}$ the total population weight. The headline metric averages Δ_e over the ten commercial street segments nearest the site. Without calibrated pedestrian counts this is an index comparison, not a volume estimate; the correlation of betweenness-type indices with observed pedestrian volumes is supported by the network-analysis literature but is context-dependent.

5.3 Consistency diagnostic

$$\text{spending per walk trip} = \frac{\sum_j C_j^{walk}}{365\,T}. \tag{9}$$

Because T counts *all* walking (exercise, transit access, dog walks) while the numerator is shopping only, a plausible value is a few dollars per average trip (implying \$5–\$15 per actual shopping trip). Values outside \$0.25–\$30 add a consistency flag to the report.

6 Employment ledger

$$J^- = \frac{\text{existing commercial sqft}}{\texttt{sqft\_per\_office\_job}}, \qquad J^+ = \frac{\text{proposed retail sqft}}{\texttt{sqft\_per\_retail\_job}}, \qquad \Delta J = J^+ - J^-. \quad (10)$$

Space-per-worker values (300 [200–450] office; 500 [400–700] retail, gross) sit between dense workplace-utilization benchmarks and the building-stock averages observable in the Commercial Buildings Energy Consumption Survey (CBECS) [19]. These are site-activity screens only; no indirect or induced employment is modeled.

7 Fiscal module

The module reports both of the two canonical fiscal-impact framings — average(per-capita) cost and marginal cost — following Burchell & Listokin's classification [10] and the planning literature's standard caution that neither is universally correct for infill [11].

7.1 Revenue

Real-estate (RE) tax lines use assessed value (AV) at the adopted rate r per \$100 of value, with p_{sqft} the assumed commercial value per square foot (`commercial_value_per_sqft`), p_{office} the assumed office value per square foot (`office_value_per_sqft`), p_{unit} the assumed assessed value per dwelling unit (`residential_value_per_unit`), and η the net-new share of on-site sales (`net_new_share`); BPOL is Virginia's Business, Professional and Occupational License tax on gross receipts.

$$\text{current RE tax} = AV_{now} \cdot r/100 \quad (11)$$

$$AV_{proj} = U \cdot p_{unit} \ + \ \text{retail sqft} \cdot p_{sqft} \ + \ \text{office sqft} \cdot p_{office} \quad (12)$$

$$\text{projected RE tax} = AV_{proj} \cdot r/100 \quad (13)$$

$$\text{RE tax increase} = (AV_{proj} - AV_{now}) \cdot r/100 \quad (14)$$

$$\text{meals tax} = S_{\text{food away}} \cdot \theta_{\text{food}} \cdot r_{\text{meals}} \quad (15)$$

$$\text{local sales tax} = (S_{\text{grocery}} + S_{\text{comparison}} + S_{\text{convenience}}) \cdot \theta_{\text{all}} \cdot r_{\text{sales}} \quad (16)$$

$$\text{personal property} \approx N_{hh} \cdot \nu_{hh} \cdot \bar{v}_{veh} \cdot r_{pp}/100 \qquad \text{(rough estimate)} \quad (17)$$

$$B_{onsite} = \text{retail sqft} \cdot g_{sf} \ - \ C_{own} \quad (18)$$

$$\text{retail BPOL} \approx B_{onsite} \cdot r_{bpol}/100 \cdot \eta \qquad \text{(rough estimate)} \quad (19)$$

$$\text{office BPOL} \approx \text{office sqft} \cdot g_{office} \cdot r_{bpol,off}/100 \qquad \text{(rough estimate)} \quad (20)$$

$$\text{BPP} \approx (\text{retail} + \text{office})\,\text{sqft} \cdot p_{bpp} \cdot r_{bpp}/100 \qquad \text{(rough estimate)} \quad (21)$$

$$\text{on-site meals tax} \approx B_{onsite} \cdot \phi_{rest} \cdot r_{\text{meals}} \cdot \eta \qquad \text{(rough estimate)} \quad (22)$$

$$\text{on-site sales tax} \approx B_{onsite} \cdot r_{\text{sales}} \cdot \eta \qquad \text{(rough estimate)} \quad (23)$$

Assessed value per unit. p_{unit} (`residential_value_per_unit`) is a first-class assumption whose default is seeded from assessment comparables of the project's own product class: the median per-unit value with 25th–75th percentile bounds, over parcels built within a 15-year lookback. Product class is selected from the project's residential tenure. For-sale projects are valued against condominium-class comparables and rental projects against apartment-class parcels (land-use code 352); an assessed condominium unit is its own account, so its per-unit value is the account total, while an apartment building holds all units on one account and per-unit value is AV_i/units_i. Applying apartment comparables to for-sale condominiums understates assessed value several-fold and is the reason tenure is extracted at all. When no comparables of the right class are available, p_{unit} falls back to a deliberately wide screening default and the report says so explicitly. Because p_{unit} is an assumption rather than internal arithmetic, it is adjustable by the reader and is ranked in the report's sensitivity list.

Displacement adjustment on on-site sales. Taxes on the project's own commercial sales are scaled by the net-new share η (`net_new_share`). A new restaurant collects meals tax, but most of its receipts are captured from existing restaurants in the same jurisdiction, which stop collecting it; only the net-new portion — new resident and visitor demand, plus sales recaptured from outside the city — adds to the city's base. Counting gross on-site receipts, as applicant fiscal analyses conventionally do, overstates the municipal gain. Equation (18) also subtracts C_{own}, the new residents' own modeled spending at the project's ground floor, which is already counted at full weight in the resident-side capture lines and is genuinely new to the city; without that subtraction it would be both double-counted and wrongly discounted. Office BPOL and business tangible property (BPP) are *not* scaled by η: professional and medical practices largely serve regional demand, and equipment is physically new to the city.

All rates (r, r_{meals}, r_{sales}, r_{bpol}, $r_{bpol,off}$, r_{bpp}) are jurisdiction configuration pinned with source URL and fiscal year; an unpinned rate degrades its metric to a stated "not computed" rather than a guess. Virginia's 1% local-option sales tax applies to food for home consumption, so grocery is correctly in the taxable base; personal services are excluded as generally non-taxable.

Personal property and BPOL revenue are *rough estimates*, and are labeled as such wherever they appear. Their tax rates are pinned city policy — r_{pp} = \$4.13 per \$100 of vehicle value and r_{bpol} = \$0.20 per \$100 of retail gross receipts, both from the city's adopted-budget Rates & Levies schedule — but the bases are assumptions: vehicles per household (ν_{hh}), average vehicle value ($\bar{v}_{veh}$), and tenant gross sales per square foot (g_{sf}) have no project-specific data behind them and carry wide bounds. When per-household personal-property budget actuals (p_{pp}) are pinned, the actuals path $N_{hh} \cdot p_{pp}$ replaces Equation (17). Every such line enters incremental recurring revenue: omitting them would systematically understate revenue for taxes the city does levy. (Car-tax relief under Virginia's Personal Property Tax Relief Act is a fixed state block grant, so marginal vehicles yield the city the full levy; the services BPOL classification at \$0.27 per \$100 is slightly above the retail rate applied here.) The office BPOL rate $r_{bpol,off}$ = \$0.40 per \$100 is the financial and professional services class, and r_{bpp} = \$4.13 per \$100 is the single personal property rate the city's schedule sets — it defines no separate business-tangible class.

Membership in the revenue roll-up is explicit: each recurring-revenue metric registers itself as it is created, and the net figures sum the registered set. An earlier version matched revenue lines against a hard-coded list of metric names, so a newly added line could be computed and displayed and then silently omitted from the net; a test now asserts that every positive recurring-revenue metric reaches the net.

External estimates and the cost-basis check. When the applicant's fiscal impact analysis or the city staff report states a fiscal estimate, it is extracted under the same verbatim-quote firewall as the project quantities and published as its own metric, attributed and marked as not computed here. These figures are never averaged into ours: the report states both ranges and whether they overlap, because a disagreement between independent methods is more informative to a reader than a blended number. A stated construction cost basis additionally drives a sanity check — new construction typically assesses at roughly 70–110% of hard cost, so a projected assessed value far outside that band is flagged as a probable wrong-product-class p_{unit} rather than published unremarked.

Benchmarks do, however, inform the screening defaults themselves. A 2026 cross-project comparison against City of Fairfax staff fiscal estimates (the staff figures have tracked outcomes more closely than literature midpoints) implied a marginal-cost share of 0.30–0.34 and multifamily student yields of 0.06–0.07; the defaults for `marginal_cost_factor` and `students_per_unit` were recalibrated partway toward those values (0.35 and 0.08/0.10 respectively), staying above the implied points deliberately. This is a considered revision of the defaults on local evidence, documented in each assumption's basis — not a fit: no external figure ever enters the computation, and per-project divergence from a staff estimate is still reported, not reconciled away.

7.2 Service costs and net impact

When the city's education transfer E and school enrollment N_{enr} are pinned, costs use the *school-split* model — the per-capita-multiplier refinement of [10]: school costs follow the project's own student estimate at the per-pupil rate, and only non-school costs are allocated per resident. The per-pupil rate is *net* of the state education revenue A the locality receives (basic aid and the education sales tax follow average daily membership, ADM, so a new student brings that revenue with them; only the local share is a cost to city taxpayers). In the formulas below, GF is the city's total General Fund expenditure and P its citywide population.

$$S_{K12} = U \cdot \texttt{students\_per\_unit} \tag{24}$$

$$\text{naive cost} = N_{res} \cdot \frac{GF - E}{P} \;+\; S_{K12} \cdot \frac{E - A}{N_{enr}} \tag{25}$$

$$\text{marginal cost} = N_{res} \cdot \frac{GF - E}{P} \cdot \mu \;+\; S_{K12} \cdot \frac{E - A}{N_{enr}}, \qquad \mu = \texttt{marginal\_cost\_factor} \tag{26}$$

$$\text{net}_{method} = R_{inc} - \text{cost}_{method} \tag{27}$$

where R_{inc} is *incremental* recurring revenue (the real-estate component is the tax *increase* when the current tax is known). The citywide identity is preserved: total allocable local cost is $(GF - E) + (E - A) = GF - A$, i.e. the state-financed share of the General Fund is excluded from local cost, exactly as a local-fiscal test requires. If A is unpinned the per-pupil rate degrades to the gross E/N_{enr}. The school term appears unscaled in both framings: pupils generate tuition-contract costs one-for-one, so a development generating few students carries proportionally low school costs instead of the school-heavy citywide average — without the split, a zero-student building would still be billed average school costs for every resident. The student generation rate is set toward the high-rise/small-unit end of the multifamily evidence [12] because university-adjacent rental buildings house predominantly single and roommate households rather than families. The naive framing remains an upper bound on the non-school side (it allocates fixed citywide costs to incremental residents); the marginal framing scales only the non-school share expected to grow and can understate costs if capacity thresholds are crossed. The headline range spans both framings; per-method nets are also published. If E or N_{enr} is unpinned, both framings degrade to the single per-capita form $N_{res} \cdot GF/P$ with the student count reported as a note.

8 Assumptions

Central values with propagated bounds; every metric lists the assumption keys that touched it.

Key	Value	Range	Basis
`occupancy_rate`	0.95	0.90–0.97	stabilized MF occupancy; cf. HVS vacancy [18]
for-sale	0.97	0.93–0.99	sold units occupied; low bound = absorption lag
`avg_hh_size_multifamily`	1.8	1.5–2.2	bedroom-mix multipliers [12]
for-sale	2.0	1.6–2.4	owner-occupied MF multipliers [12]
`income_premium_new_construction`	1.125	1.0–1.25	screening premium, stated openly
for-sale	1.35	1.15–1.60	buyers anchor to price, not rent
`ces_scale`	1.0	0.85–1.15	CES national-average drift allowance [13]
`beta_walk` (min^{-1})	0.10	0.05–0.15	purpose-specific walk decay 0.09–0.11 [15]; national NHTS du
β_d (min^{-1}, fixed)	0.15	—	drive impedance decay
`walk_share_neighborhood`	0.60	0.40–0.80	zero-impedance walk preference
`walk_share_comparison`	0.10	0.00–0.30	comparison goods travel by car
`walk_share_grocery_entertainment`	0.30	0.10–0.50	interpolated
`walk_trips_per_resident_day`	1.0	0.5–2.0	NHTS metro rates + walkability uplift [14, 20]
`sqft_per_office_job`	350	250–500	CBECS gross sqft/worker; older part-vacant stock [19]
`sqft_per_retail_job`	500	400–700	planning standard; cf. CBECS [19]
`own_retail_sqft_per_equiv_poi`	2,000	1,500–3,000	typical inline suite size
`residential_value_per_unit`	comps	comp 25th–75th	median per-unit AV, product-class comps
fallback, rental	\$360k	\$275k–\$450k	screening default when no comps
fallback, for-sale	\$750k	\$450k–\$1.2M	screening default when no comps
`commercial_value_per_sqft`	\$275	\$200–\$400	screening; replace with comps
`office_value_per_sqft`	\$300	\$200–\$450	new-construction NoVA office AV (estimate)
`students_per_unit`	0.08	0.05–0.12	high-rise MF [12], shaded to local staff-estimate benchmarks
for-sale	0.10	0.05–0.16	owner-occupied condo yield [12], shaded likewise
`marginal_cost_factor`	0.35	0.25–0.55	share of average cost that scales; calibrated to local staff esti
`vehicles_per_household`	1.4	1.0–1.8	ACS vehicle norms, MF renters (estimate)
`avg_vehicle_assessed_value`	\$14,000	\$9k–\$20k	Northern Virginia vehicle rolls (estimate)
`retail_sales_per_sqft`	\$400	\$250–\$600	neighborhood retail gross sales (estimate)
`office_receipts_per_sqft`	\$500	\$300–\$800	professional-services receipts/sqft (estimate)
`bpp_value_per_sqft`	\$8	\$4–\$15	business FF&E per commercial sqft (estimate)
`net_new_share`	0.30	0.15–0.50	share of on-site sales new to the city; retail displacement
`onsite_restaurant_share_of_retail`	0.50	0.30–0.70	ground-floor tenant mix

Structural constants: walk speed 4.8 km/h; DBSCAN $\varepsilon = 150\,\mathrm{m}$, minPts = 3; network extent = city boundary + 3 km; POI study area = boundary + 2 mi; node population radius 400 m; baseline sampling $k = 2000$, fixed seed; comp lookback 15 years.

9 Implementation decisions

- **Per-business destinations.** The Huff run treats every retail POI as an individual destination (a single-source Dijkstra makes per-POI times free); DBSCAN clusters exist only as reporting rollups. Capture is therefore spatially resolved to business locations, and map heat layers consume real per-business weights rather than interpolated cluster centroids.
- **Joint rather than nested choice.** Equation (5) is a single multinomial logit over (destination, mode) pairs. A nested logit would relax the independence-of-irrelevant-alternatives (IIA) property across modes; the joint form was chosen for transparency and because screening bounds on w_c and β_w dominate the error budget.

- **Exact marginal flows.** Project-caused foot traffic is computed as the shortest-path routing of the residents' own trips, replacing an earlier sampled before/after betweenness difference whose noise exceeded the signal at this project scale. Sampling now appears only in the baseline denominator.
- **Census-block population.** Node population weights use 2020 census block centroids (the finest published geography) fetched by spatial envelope, which also extends the baseline beyond the city line that the county-scoped ACS layer stopped at.
- **Determinism.** All sampling is seeded; two evaluations on a warm cache produce identical metric values.
- **Large-language-model (LLM) firewall.** Document extraction demands a verbatim $\leq$15-word quote per numeric field, verified in code against the document text; unverifiable values are nulled. Parcel identifiers are verified the same way.
- **Narrative validator.** The generated report is scanned for numbers; each must match a metric value/bound, spec quantity, assumption, or deterministic note within rounding, else the draft is regenerated once and then hard-fails.
- **Assessment data.** The city publishes no bulk assessment layer, so site values and multifamily comparables come from the public Patriot WebPro search interface, throttled to $\sim$1 request/s and cached; owner names are neither parsed nor stored.
- **Map scales.** Heat layers are clipped to commercial-retail zoning polygons and colored on dollar-pinned viridis ramps (nice-number ceilings), so a color means the same dollars regardless of layer toggling; capture and walk-in capture share one scale to make their comparison meaningful.
- **Interactive assumption adjustment.** Each metric publishes an exact power-law decomposition over the assumptions it depends on multiplicatively (or by a fixed power, e.g. the Engel elasticities): $v' = \sum_t v_t \prod_k (a_k/a_k^0)^{e_{t,k}}$. Because the joint-choice probabilities are independent of the demand-side assumptions, this client-side recompute reproduces the pipeline's own arithmetic exactly; the evaluation asserts $\sum_t v_t = v$ on every run. Network-model parameters (β_w, mode shares, own-retail attractiveness) appear in no term and are presented as requiring a full re-run.

10 Known limitations and audit notes

Four limitations identified in the July 2026 audit have been addressed in this version; their adopted resolutions and residual caveats:

1. **Household size (addressed).** The unit-mix-blind ACS renter average ($\approx$ 3.1) was replaced by a multifamily bedroom-mix multiplier, $\bar{s}$ = 1.8 [1.5–2.2] [12]. *Residual:* the multiplier assumes a typical studio/one-bedroom/two-bedroom mix; when a project's actual bedroom program is extracted it should set $\bar{s}$ directly.
2. **Aggregate income (addressed).** Tract income now uses the mean (B19025/B11001) rather than median × households, with a median fallback under suppression. *Residual:* tract means still smooth over building-level differences; the new-construction premium carries that uncertainty.

3. **CES income scaling (addressed).** Category spending scales by r^{ε_c} per Equation (4) instead of linearly. *Residual:* the elasticities are national approximations read from CES quintile gradients, not locally estimated.

4. **Walk-trip rate (addressed).** τ reduced from 2.0 to 1.0 [0.5–2.0], anchored to NHTS metro person-level walking rates with a walkability uplift consistent with the built-environment literature [14, 20]. The foot-traffic *percentage* is unaffected (τ cancels) and walk-in dollars are independent of τ; trips and street-flow magnitudes scale with it. *Residual:* local counts remain the real calibration.

5. **IIA.** The joint logit imposes proportional substitution across all (destination, mode) pairs; a nested structure would be more standard.

6. **Transferability of β_w.** The central value (0.10) sits between the Twin Cities purpose-specific estimates (shopping 0.094, restaurant 0.093 [15]) and the national NHTS duration decay (0.073 [16]). Both are fit to observed trip-length distributions, which conflate willingness with spatial structure [17]; in a competing-destinations model the pure impedance coefficient is plausibly gentler still. Pedestrian- environment quality is not separately modeled — for hostile walking environments, read the steep end of the band (0.15).

7. **Attractiveness proxy.** Competitor attractiveness is POI count (floor area being unavailable in open POI data), while own-retail attractiveness is derived from square footage; the conversion constant carries the reconciliation and its bounds.

8. **Sales-tax base share.** The taxable base applies the all-retail in-city share to the three taxable categories; a category-specific share would be marginally more precise.

9. **Foot-traffic index is uncalibrated.** Percent changes compare modeled flows to a sampled opportunity index, not to counted pedestrians; local counts would convert the layer to volumes.

10. **Fiscal method dependence.** The net range spans average-cost and marginal-cost framings on purpose; the midpoint of that range is a presentation convenience, not an estimate with independent standing.

11. **Rate-based revenue estimates.** Personal property and BPOL revenue use pinned city rates but assumed bases, per Equations (17) and (19); they are labeled rough estimates and carry the widest relative bounds on the revenue side. Per-household personal-property budget actuals and tenant-level receipts would replace them.

11 Data sources

City of Fairfax GeoHub ArcGIS services (boundary, parcels, zoning); City of Fairfax public assessment database (Patriot WebPro); City budget pages (tax rates, General Fund, pinned with fiscal year) and the adopted-budget Rates & Levies schedule (personal property and BPOL rates); City of Fairfax Schools fiscal-year 2026 budget announcement (tuition contract with Fairfax County Public Schools and projected ADM, for the school-split cost model); Census ACS 5-year (B01003, B19013, B25010, B25044, B08301) and 2020 PL 94-171 block population via TIGERweb; Census Longitudinal Employer–Household Dynamics Origin–Destination Employment Statistics (LEHD LODES) version 8; OpenStreetMap via OSMnx [9]; Overture Maps places; BLS Consumer Expenditure Survey 2023 [13]; FHWA National Household Travel Survey [14]; Virginia Code §58.1-605 (local-option sales tax).

References

[1] Huff, D. L. (1963). A probabilistic analysis of shopping center trade areas. *Land Economics*, 39(1), 81–90.

[2] Huff, D. L. (1964). Defining and estimating a trading area. *Journal of Marketing*, 28(3), 34–38.

[3] Wilson, A. G. (1971). A family of spatial interaction models, and associated developments. *Environment and Planning A*, 3(1), 1–32.

[4] McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics* (pp. 105–142). Academic Press.

[5] Ben-Akiva, M., & Lerman, S. R. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press.

[6] Freeman, L. C. (1977). A set of measures of centrality based on betweenness. *Sociometry*, 40(1), 35–41.

[7] Brandes, U. (2001). A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2), 163–177.

[8] Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. *Proceedings of KDD-96*, 226–231.

[9] Boeing, G. (2017). OSMnx: A Python package to work with graph-theoretic OpenStreetMap street networks. *Journal of Open Source Software*, 2(12), 215.

[10] Burchell, R. W., & Listokin, D. (1978). *The Fiscal Impact Handbook: Estimating Local Costs and Revenues of Land Development*. Center for Urban Policy Research, Rutgers University.

[11] Kotval, Z., & Mullin, J. (2006). *Fiscal Impact Analysis: Methods, Cases, and Intellectual Debate*. Lincoln Institute of Land Policy working paper WP06ZK2.

[12] Listokin, D., Voicu, I., Dolphin, W., & Camp, M. (2006). *Who Lives in New Jersey Housing? Updated New Jersey Demographic Multipliers*. Center for Urban Policy Research, Rutgers University.

[13] U.S. Bureau of Labor Statistics. *Consumer Expenditure Surveys*, 2023 tables. `https://www.bls.gov/cex/`

[14] Federal Highway Administration. *National Household Travel Survey*, 2022. `https://nhts.ornl.gov/`

[15] Iacono, M., Krizek, K., & El-Geneidy, A. (2008). *Access to Destinations: How Close Is Close Enough? Estimating Accurate Distance Decay Functions for Multiple Modes and Different Purposes*. Report MN/RC 2008-11, Minnesota Department of Transportation. Table 10 estimates negative-exponential walk decay by travel time: work 0.106, shopping 0.094, restaurant 0.093, recreation 0.100 min^{-1}. Journal version: Iacono, Krizek & El-Geneidy (2010), *Journal of Transport Geography*, 18(1), 133–140.

[16] Yang, Y., & Diez-Roux, A. V. (2012). Walking distance by trip purpose and population subgroups. *American Journal of Preventive Medicine*, 43(1), 11–19. National 2009 NHTS; fitted duration decay 0.073 min^{-1} ($R^2 = 0.99$); median walk trip 10 minutes; 65% of walk trips exceed one quarter mile.

[17] Sheppard, E. S. (1995). Modeling and predicting aggregate flows. In S. Hanson (Ed.), *The Geography of Urban Transportation* (2nd ed., pp. 100–128). Guilford Press.

[18] U.S. Census Bureau. *Housing Vacancies and Homeownership (CPS/HVS)*. `https://www.census.gov/housing/hvs/`

[19] U.S. Energy Information Administration. *Commercial Buildings Energy Consumption Survey (CBECS)*, 2018 building characteristics tables. `https://www.eia.gov/consumption/commercial/`

[20] Ewing, R., & Cervero, R. (2010). Travel and the built environment: A meta-analysis. *Journal of the American Planning Association*, 76(3), 265–294.