Houston Land Value Tax: a parcel-level analysis

What if the City of Houston taxed land instead of buildings?

Houston LVT working draft

2026-04-25

Abstract

We estimate what would happen if the City of Houston replaced its current property tax with a revenue-neutral land value tax (LVT) on the City's tax year 2025 (TY2025) appraisal roll. At the City's adopted FY2026 rate of $0.5192 per $100 of taxable value, full-market property tax revenue from the 655,110 taxable parcels inside the City of Houston comes to $2.06 billion. Solving for revenue neutrality on the same parcels yields an LVT rate of **$1.31 per $100 of land value** — about 2.5× the current rate. Under that switch, **53% of taxable parcels would pay less; 42% would pay more**. The bill on a median single-family home falls by 15.7% (≈ $217); the bill on a typical in-city vacant lot rises by roughly 150% (≈ $400 to $1,000+). Surface lots and underutilized commercial parcels carry most of the new burden; high-rise condos and dense multifamily save the most. We use HCAD's reported land values directly (Texas's non-disclosure laws prevent a comp-based market model on public data), which makes the per-parcel results reported here *conservative*: the real LVT gap on surface lots and vacant land is almost certainly larger than what we show.

Executive summary

Solved revenue-neutral LVT rate	**$1.3098 / $100 of land value**
Multiple of current rate	2.52×
Taxable parcels in City of Houston	655,110
Off-roll (state-class X*)	30,303
Total uncapped current tax	$2.058 B
Solved-for total LVT revenue	$2.058 B (by construction)
Total land market value	$157.1 B
Total improvement market value	$239.3 B
Land share of taxable market value	**39.6 %**

Savers under LVT	349,383 (53.3 %)
Payers under LVT	274,754 (41.9 %)
Roughly unchanged (|Δ%| < 5 %)	30,973 (4.7 %)

One-line policy story. A revenue-neutral land value tax would shift City of Houston property tax burden off improvements (homes, offices, factories) and onto land — most visibly onto surface parking lots and vacant in-city land, which today pay almost nothing relative to the value of the land they occupy.

This is a working draft. The map at https://powerfundamental.com shows the parcel-by-parcel results visually; this paper documents the methodology and limitations so readers can decide how much weight the headline numbers deserve.

---

1. Why study LVT for Houston

A land value tax bills you on the assessed value of your land but exempts any buildings, paving, and other improvements on it. The economic argument for LVT rests on a few standard claims:

1. **Land supply is fixed**, so a tax on land cannot be passed on to renters or shifted to consumers in the way taxes on labor, capital, or buildings can. The incidence falls on the landowner.
2. **Owners of underused land are penalized for holding it idle**, which creates a marginal incentive to sell, build, or improve. Surface parking and vacant lots in dense areas are the canonical examples.
3. **Improvements are no longer taxed**, which removes a small but real disincentive to renovate, expand, or build new housing.

Houston is an unusually instructive city to model. It has:

- No traditional zoning (most of the country's largest no-zoning city), so "what would happen if owners could freely build more" is less hypothetical than it would be in, say, San Francisco.
- A large stock of **surface parking** in the urban core — Downtown, Midtown, the Texas Medical Center, and the freeway frontage of every loop. Surface lots are the parcels for which LVT and current property tax diverge the most.
- A property-tax system loaded with **caps and exemptions** (homestead, over-65, disabled vet) that obscure the underlying market values. Stripping those caps to compare "full-market property tax" against "full-market LVT" isolates the policy effect of switching the *base* of the tax from "land + building" to "land only".
- A recent (TY2025) certified appraisal roll plus a published FY2026 tax rate, so we can anchor revenue neutrality against numbers the City itself publishes.

The point of the analysis is not to recommend that Houston adopt LVT — that question is partly empirical and partly political. It is to make the *distributional* consequences visible at the parcel level:

who pays more, who pays less, and by how much, on real Houston tax-roll data.

2. Scope and methodology decisions (locked in)

These were settled before the analysis ran and are documented here for transparency:

- **Geographic scope.** City of Houston only — a centroid-in-polygon filter against the official COH boundary published by the City's GIS division. We exclude unincorporated Harris County and other Harris County cities (Pasadena, Pearland, Sugar Land, etc.), even when they share appraisal-district infrastructure with Houston.
- **Tax jurisdiction.** City of Houston tax only. HISD, Harris County, MUDs, and other overlapping taxing units are out of scope; they have separate funding constraints and aren't policy levers the City controls.
- **LVT structure.** Pure LVT. No improvement tax, no progressive bands. We solve for the single uniform rate that yields the same total revenue from the same set of taxable parcels.
- **Caps and exemptions.** Stripped on **both sides** of the comparison. We apply no homestead exemption, no over-65 freeze, no disabled-vet exemption. HCAD publishes both market and capped values; we use market consistently. This isolates the LVT-vs.-property-tax effect from the separate question of whether to keep caps and exemptions. The capped/exempt-applied comparison is interesting but blends two policy effects.
- **Tax-exempt parcels.** Parcels whose HCAD `state_class` begins with "X" (X1 government, X2 charity, X3 religious, etc.) are excluded from both the revenue total and the per-parcel comparison. They are tracked separately as "off-roll" — see *Limitations*.
- **Land valuation source.** HCAD's reported `land_val` for every parcel. The original plan called for a comp-based hedonic regression on HCAD sales; Texas's status as a non-disclosure state makes this infeasible on public data (see §4).

3. Data sources

All data is public and free.

Source	Used for	Vintage
HCAD `Parcels.zip` (parcel polygons)	Geographic basis	2026-04-15 release

Source	Used for	Vintage
HCAD `real_acct.txt`	Per-parcel land value, building value, state class, neighborhood	TY2025 certified
HCAD `land.txt`	Per-parcel HCAD land-use code and supporting detail	TY2025 certified
HCAD `Code_description_real.zip`	State-class and land-use code dictionaries	TY2025
City of Houston Open Data Portal	Official COH boundary, super-neighborhoods, TIRZ, METRO Rail, Walkable Places, TOD, Major Thoroughfare	Latest
OpenStreetMap (Overpass API)	`amenity=parking` polygons	2026-04
Microsoft USBuildingFootprints (Texas)	Building footprints for parking-area subtraction	2019–2020 ML extraction
FEMA NFHL ArcGIS REST	Flood hazard zones (informational; not used in tax math)	Current
TY25 Notice of Meeting to Vote on Tax Rate (COH)	Adopted FY2026 rate, certified TY2024→TY2025 levy	Adopted Oct 15 2025

The pipeline is fully reproducible from public sources; see §8.

For TY2025, HCAD published 1,538,086 parcel polygons across all of Harris County. After centroid-in-polygon filtering against the COH boundary we keep **685,413 parcels**. Of those, 681,711 (99.5 %) have a matching account-level record in `real_acct.txt` ; the remaining 3,702 are mostly stacked-condo fragments and special-purpose accounts that don't carry an attribute record on their own.

4. Land valuation: the elephant in the room

A market-based parcel-level valuation model is the load-bearing technical step in any LVT analysis worth taking seriously. Our original plan — outlined in the project repository's `project_plan.md` — was a hedonic regression of land price as a function of parcel size, location (distance to downtown, METRORail stations, employment centers), submarket (super-neighborhood, target-encoded), policy overlays (TIRZ, Walkable Places, FEMA flood zones), shape (perimeter²/area compactness, corner-lot flag), and recent-sales adjacency, trained on land sales from HCAD's published `deeds.txt` and applied to every taxable parcel.

That plan does not survive contact with the data. **Texas is one of a small number of non-disclosure states**: parties to a real-estate transfer are not required to publish or report the price, and HCAD's public files do not include sale prices anywhere. `deeds.txt` records account number, date of sale, and clerk reference — but no consideration. Without sale prices we cannot fit a hedonic model on public data alone.

We considered three responses:

1. **Buy commercial sale data.** ATTOM, CoreLogic, and similar vendors maintain Texas sale records reconstructed from clerk consideration fields, private MLS data, and proprietary methods. Cost: ~$5,000/year for a Harris County dataset, plus license restrictions on republishing. Not done.
2. **Apply a uniform multiplier from Texas Comptroller PTAD ratio studies.** The Comptroller publishes biennial sales-ratio audits per appraisal district. These let us scale HCAD's values toward market in the aggregate, anchored to the State's own audit methodology. Worthwhile follow-up; not done in v1.
3. **Use HCAD's reported land values directly.** Defensible, transparent, and what this paper does.

The mathematical consequence is that **our parcel-by-parcel LVT comparison is conservative on its headline finding**. HCAD's land valuations are widely understood to be biased low relative to market, particularly for vacant and underutilized parcels in appreciating submarkets — the agency's protest-driven defensible-appraisal incentive structure pushes against aggressive land-value updates. To the extent that HCAD systematically underestimates land relative to improvements, the ratio Σ land_val / Σ tot_mkt_val we compute is too low, the solved LVT rate we report is too high, and the implied burden shift onto underutilized land is too small. The *real-world* gap between current property tax and a revenue-neutral LVT, on the kinds of parcels we identify as paying disproportionately more, is almost certainly larger than what this paper shows.

5. Tax-comparison methodology

For each taxable parcel (HCAD account inside the COH boundary, state class not starting with `X`):

```
improvement_market_value = bld_val + x_features_val
total_market_value       = land_val + improvement_market_value
current_tax_uncapped     = total_market_value × current_rate
```

with `current_rate = 0.0051919` (i.e. $0.5192 per $100), the City's adopted FY2026 rate.

Aggregating over the 655,110 taxable parcels:

```
total_current_revenue = Σ current_tax_uncapped  = $2,057,949,599
total_land_value      = Σ land_val              =   $157,124,671,507
```

```
lvt_rate              = total_current_revenue / total_land_value
                      = 0.013098 = $1.3098 / $100
```

Per parcel:

```
lvt_tax   = land_val × lvt_rate
Δ         = lvt_tax − current_tax_uncapped
Δ%        = Δ / current_tax_uncapped
```

We tag each taxable parcel as a *saver* (Δ < 0 and |Δ%| ≥ 5 %), *payer* (Δ > 0 and Δ% ≥ 5 %), or *neutral* (|Δ%| < 5 %).

6. Results

6.1 Aggregate

The solved LVT rate of **$1.3098 / $100** is 2.52× the current rate. This is *below* the 4–8× rule-of-thumb band that LVT studies of US cities typically report. The reason is mechanical: that rule of thumb assumes land is 15–25 % of total taxable value, and HCAD's appraisals put Houston's land share at **39.6 %**.

There are two ways to read the high land share:

1. **Houston really does have a high land share** — rapid in-city land appreciation, a large surface-parking footprint, and lots of low-density peripheral development genuinely push the ratio above the 25 % typical of a denser US city.
2. **HCAD undervalues improvements** relative to land. That would mechanically inflate the land share. We can't tell the two apart on public data alone; sale-data work would.

Both readings are consistent with the qualitative finding (LVT shifts burden off improvements onto land) but they pull the rate multiple in opposite directions. The white-paper headline rate of 2.5× should be read as "approximately the right order of magnitude given HCAD's appraisals" rather than a sharp prediction.

The **uncapped revenue total of $2.058 B is 1.45× the certified FY2026 levy of $1.421 B**. The gap ($637 M) is the implicit revenue subsidy from homestead caps, over-65 freezes, and exemption claims on the current property-tax base. It is itself a publishable headline: ending caps and exemptions on the current base, holding everything else constant, would raise roughly $640 M/year in additional City of Houston property tax revenue.

6.2 By use category

Aggregating taxable parcels by the first letter of HCAD's state class:

Use bucket	Parcels	Current tax	LVT tax	Δ total	Median Δ%
Single-family residential (A*)	476,251	$1,130 M	$1,227 M	+$97 M	**–15.2 %**
Multifamily (B*)	12,075	$253 M	$155 M	**–$98 M**	–7.3 %
Vacant (C*)	63,561	$45 M	$113 M	**+$68 M**	**+152 %**
Commercial / industrial (F*)	36,246	$558 M	$512 M	–$46 M	+30.5 %
Condo (Z*)	58,236	$62 M	$30 M	–$32 M	–52.1 %
Other (D, *E*, J*, etc.)	8,741	$10 M	$20 M	+$10 M	+152 %

The headline distributional finding is in the second-to-last column of the *median* row — the typical parcel of each kind:

- **Median single-family home pays 15.2 % less under LVT.** On the median Houston A1 parcel, that's a drop from $1,479 to $1,113 of City of Houston property tax — a **$217 annual saving**.
- **Median vacant lot (C1) pays 152 % more under LVT** — a rise from $262 to $661, a $399 annual increase. Vacant commercial (C2) sees the same multiplier on a higher base: $553 → $1,396, a $843 increase. C3 (vacant not in-city) is a small bucket and behaves the same way.
- **Median commercial (F1) parcel pays 30.5 % more under LVT.** On the median F1 that's $3,848 → $4,775, a $615 annual increase. The total category *drops* $46 M under LVT because a few large-improvement properties (office towers, malls) save substantially more than the median.
- **Median industrial (F2) parcel pays 21 % more under LVT** despite a larger median property: industrial parcels are typically improvement-heavy on cheaper land, so a 21 % median increase on a high base is consistent with the story.
- **Multifamily and condo parcels save the most.** Both are dense improvement-on-land — the policy story is that LVT rewards intensive use. The condo result deserves a footnote: condo land values are split among many unit owners by HCAD's "Z" coding, so each individual unit sees only a sliver of land tax. Aggregated as a category, condo parcels pay 52 % less under LVT (median); the practical effect is concentrated in dense central multifamily.

6.3 Surface parking and underutilized land — the flagship finding

The parcels that pay the most more under LVT are the ones that hold valuable land without putting much improvement on it. Vacant parcels (the C* row above) pay 152 % more on the median, but that bucket understates the issue because the centroid-driven HCAD valuations aren't even tracking the *full* market value of land in walkable in-city neighborhoods. On the parcel inventory we build for the visual map (project plan §3), three categories of underutilized land emerge:

- **Vacant lots and tracts** (HCAD state class C1/C2/C3/XV, or zero improvement value with state-class not X/J): **74,814 parcels**.
- **Parking-primary parcels** — host parcels where the surface-parking coverage (OSM + footprint subtraction) is greater than 50 % of the parcel area: **18,773 parcels**.
- **Low-improvement parcels** — non-vacant parcels where improvement-value/land-value < 0.20: **32,759 parcels** (some overlap with the parking-primary set).

These three categories sum to roughly 126,000 parcels — about 19 % of taxable parcels in the City of Houston. They are exactly the parcels for which LVT and current property tax diverge most in dollar terms. The interactive map visualizes them as a separate toggleable layer.

6.4 Cross-check against LVTShift's distributional framework

The LVTShift toolkit (Center for Land Economics) maintains a parallel modeling pipeline for 13+ US cities, none of which are in Texas. Their `validate.md` skill specifies expected directional signs of tax change by property category under a standard 4:1 split-rate scenario, used as a sanity gate when adding a new city. We applied their 4:1 split-rate model to our Houston dataset (via `analysis/cross_check_lvtshift.py`) and check the result against their gate:

Category	Our median Δ% (4:1 split)	LVTShift expected band	Pass?
Vacant (HCAD C*)	+82.7 %	+50 % to +250 %	✓
Single-family (A*)	−8.2 %	−30 % to −3 %	✓
Multifamily (B*)	−4.0 %	−50 % to −1 %	✓
Commercial / industrial (F*)	+16.6 %	−50 % to +60 %	✓
Condo (Z*)	−28.3 %	−80 % to 0 %	✓

All five categories with explicit expectations land in LVTShift's predicted band — a useful soft validation that our data ingestion and tax math are not producing nonsense relative to an independent reference implementation, despite using a different appraisal authority, different state non-disclosure regime, and different scenario primitive (pure LVT vs split-rate).

For completeness, applying their split-rate solver across a range of land:improvement ratios on our Houston revenue target gives:

Scenario	Land rate ($/$100)	Improvement rate	Multiple of current
Current (uniform)	0.5192	0.5192	1.00×
Split-rate 2:1	0.7436	0.3718	1.43×
Split-rate 4:1	0.9486	0.2372	1.83×
Split-rate 10:1	1.1367	0.1137	2.19×
Pure LVT (this paper)	**1.3098**	**0**	**2.52×**

Pure LVT is the limit of split-rate as the ratio goes to infinity. Houston's high land share (39.6 %) means even a moderate 2:1 split already shifts a meaningful amount of burden onto land; the curve continues out to the limit case of pure LVT, with diminishing distributional change at higher ratios.

Cross-check artifact: `data/final/lvtshift_crosscheck.json` .

6.5 Off-roll land — the back-filled estimate

HCAD does not, as a rule, appraise tax-exempt property carefully. The **25,553** state-class-X parcels inside the City of Houston (Texas Medical Center, downtown government complex, universities, churches, public schools, parks, major-roadway ROW) carry a HCAD-reported total land value of just **$0.24 B** — an order of magnitude too small. The agency reports `land_val = 0` for most of these because there's no statutory tax obligation, not because the land is worthless.

To produce a defensible floor estimate, we apply the LVTShift-style imputation pattern: for each super-neighborhood, compute the median land $/sqft from taxable parcels, then apply that median by area to every exempt parcel in the same neighborhood. Parcels outside any super-neighborhood (annexation slivers, far-edge tracts) fall back to the citywide median of ~$11/sqft. Code: `analysis/exempt_backfill.py` .

HCAD-reported off-roll land value	**$0.24 B**
Synthetic off-roll land value (floor)	**$70.0 B**
Implicit ratio (synthetic ÷ HCAD)	**288 ×**
Off-roll share of *total* Houston land value	**30.8 %**
Implied revenue if taxed at current rate ($0.519/$100)	**$363 M / yr**
Implied revenue if taxed at LVT rate ($1.31/$100)	**$917 M / yr**

The 30.8 % figure is the headline. Roughly **one-third** of all the *land value* in the City of Houston, by our floor estimate, sits off the property tax rolls — a fraction that any policy conversation about taxing land needs to start with. (For comparison, the LVT rate-solver in §6.1 treats off-roll parcels as off-roll on *both* the current and LVT sides, preserving the apples-to-apples comparison; the off-roll figure is a *separate* policy headline about base erosion, not a critique of the LVT math itself.)

The largest individual exempt parcels by floor synthetic value:

State class	Site address	Acres	Synthetic land value
X1	1001 E Memorial Loop Dr (Memorial Park)	1,392	$5.15 B
X1	0 Katy Fwy (highway ROW)	2,742	$1.37 B
X1	0 Highway 6 S (highway ROW)	1,819	$0.91 B
X1	6100 Main St (TMC area)	185	$0.83 B
X1	0 Beeler Rd	1,253	$0.63 B

The TMC entry deserves a flag: 6100 Main St is one parcel inside the Texas Medical Center complex, and our floor estimate of $0.83 B is itself almost certainly low. TMC's *true* land value is plausibly an order of magnitude higher — its 1,300+ acres of contiguous land in a globally important medical-research cluster, half a mile from downtown, is not adequately captured by the median-$/sqft of nearby small commercial parcels. A v2 might apply a centrality multiplier to the largest exempt parcels; v1 reports the floor.

Output artifacts: `data/interim/parcels_exempt_synthetic.parquet` (per-parcel) and `data/final/exempt_backfill_summary.json` (aggregate JSON).

6.6 Equity by income and demographic composition

For every taxable parcel we spatial-join its centroid to its 2022 ACS 5-year block group, then bucket parcels into quintiles by the block group's median household income and by minority share (1 – non-Hispanic white share). Coverage: 99.99 % of parcels match a block group; 88 % of parcels' block groups have a non-suppressed median-income estimate.

Results below are the median Δ% (LVT tax / current-tax-uncapped – 1) of the parcels in each quintile, plus the share that come out as savers.

By block-group median household income (taxable parcels):

Quintile	Income range	Parcels	Median Δ%	Savers
Q1 lowest	$9 k – $44 k	121,541	**+5.6 %**	47.7 %
Q2	$44 k – $61 k	121,577	–17.0 %	59.4 %
Q3	$61 k – $82 k	121,378	**–29.6 %**	**68.8 %**
Q4	$82 k – $137 k	121,873	–13.9 %	57.4 %
Q5 highest	$137 k – $250 k	121,058	+2.2 %	48.8 %

By block-group minority share (taxable parcels):

Quintile	Minority share	Parcels	Median Δ%	Savers
Q1 most-white	0 % – 39 %	131,432	+0.9 %	49.5 %
Q2	39 % – 68 %	130,991	–18.8 %	60.2 %
Q3	68 % – 90 %	130,806	**–22.5 %**	**61.6 %**
Q4	90 % – 97 %	130,756	–15.2 %	58.0 %
Q5 most-diverse	97 % – 100 %	130,863	+0.2 %	49.9 %

Reading: the **middle of both distributions wins** under LVT. Houston's working-class and middle-income neighborhoods, and Houston's mixed and majority-minority neighborhoods, save the most. The two extremes — the lowest-income tail, and the most-segregated white tail — come out roughly neutral on the median.

Two caveats deserve to be stated alongside these numbers:

1. **The lowest-income quintile pays slightly more (+5.6 %), not less.** This is a known LVT policy property, not a bug. In neighborhoods where small or aging buildings sit on relatively valuable land — the same "low-improvement" pattern that drives the surface-lot story — LVT shifts burden onto the land. Some of those parcels are owned by lower-income homeowners. A pure LVT, applied with no homestead protection (which we strip on both sides for the apples-to-apples comparison), exposes them. In a real implementation, a circuit breaker or a homestead exemption on the LVT side would mitigate this; that's a policy lever this paper doesn't model.

2. **These are *block-group* attributes, not parcel-level demographics.** We're measuring "parcels in lower-income neighborhoods" and "parcels in majority-minority neighborhoods", not the income or race of the parcel *owner*. The owner-level distribution would require linking to an ownership-demographics dataset that isn't in HCAD's public release.

Output artifacts: `data/interim/parcels_with_demographics.parquet` (per-parcel) and `data/final/equity_summary.json` (quintile aggregates).

6.7 Caps and exemptions

The single largest single-axis comparison the analysis surfaces is not the current-vs.-LVT split itself but the **uncapped-vs.-capped** gap on the current base:

	Amount
Σ current-tax-uncapped on taxable parcels (full market)	$2.058 B
City's certified FY2026 levy (capped + exempt-applied)	$1.421 B
Implicit subsidy from caps + exemptions	**$637 M / year**

This number is a side-effect of stripping caps to compare apples to apples; it deserves separate billing in the policy conversation. The protective effect of the homestead cap, the over-65 freeze, and the disabled-veteran exemption on Houston's most-protected residential parcels currently shifts roughly $640 M of would-be City of Houston property tax obligation off long-tenured residents and onto everyone else (renters, newer owners, commercial property, the unprotected). Whether that protection is good policy is a separate question; surfacing the dollar magnitude is the contribution here.

7. Limitations and assumptions

1. **HCAD land values, not market.** §4 covers this in detail. Headline: results are conservative on the surface-lot/vacant-land delta. The real LVT gap on these parcels is almost certainly larger.

2. **Off-roll parcel valuations are back-filled by neighborhood median.** HCAD reports `land_val = bld_val = 0` for most tax-exempt parcels — the agency has no statutory incentive to appraise property that won't be taxed. We synthesize a floor estimate of off-roll land value by computing the median $/sqft of land from taxable parcels in each super-neighborhood and applying it to each exempt parcel by area; see §6.5 for the headline number. Caveat: the median-based recipe under-estimates uniquely large exempt parcels (Texas Medical Center, downtown government complex, universities) because their improvements have specialty value the neighborhood median can't capture. The figures we report are floors.

3. **Stacked-condo polygons inflate parcel-area sums but not tax math.** HCAD's parcel layer includes one polygon per condo unit; 64,633 of the 685,413 parcels in our set are flagged `Stacked = 1`. Their polygons share a ground footprint with the building's other units, so summing parcel areas produces a city-area total of 1,095 sq mi against an actual city extent of 645 sq mi. Each parcel still carries its own tax obligation, so the LVT math is unaffected. Anyone deriving "total acres of X" from the parcel layer should restrict to `Stacked = 0` ground parcels.

4. **Centroid filter has 0.2 % outside-boundary leakage.** Two-tenths of one percent of the kept parcel area technically extends past the COH boundary (large peripheral parcels straddling the city limit). Negligible for tax math; mentioned for completeness.

5. **Hand-validation precision (200 of 300 parcels walked through, 2026-04-26).** A stratified sample (100 each of vacant, parking-primary, low-improvement) was checked against satellite imagery; the low-improvement bucket has not been validated yet. Headline precision:

Bucket	Precision	n judged
parking-primary	**97 %** (96/99)	99 (1 unsure)
vacant	**78 %** (70/90)	90 (10 unsure)
low-improvement	not yet validated	—

   The vacant rule splits cleanly into two sub-rules:

   - `state_class ∈ {C1, C2, C3, XV}` → **92 %** precision (48/52). Solid.
   - `bld_val + x_features_val == 0` → **58 %** precision (22/38). Noisy.

   The noisy half was sweeping in X-class parcels (government, charity, religious — physically built, just exempt) and J-class parcels (utilities, railroads — physically present infrastructure). HCAD reports `bld_val = 0` for these because it doesn't bother appraising tax-exempt or non-improvable property, not because the parcels are physically vacant. Pipeline patch on 2026-04-26 excludes state-class first-letters X and J from the zero-improvement vacant rule (`pipeline/05_classify.py`). After the patch, vacant count drops from 102,138 to 74,814; the 27,324 reclassified parcels move to "other" (still off-roll if X-class, still taxable but not bucketed otherwise). The validated precision should land near the 92 % of the C/XV-rule. The validation export is checked in at `analysis/validation/results/validation_2026-04-26.json` .

6. **Layer B parking polygons are an upper-bound mask.** Our parking-detection pipeline produces 54,948 polygons across two layers: ~20k from `amenity=parking` in OpenStreetMap, ~35k from a footprint-subtraction procedure that takes (parcel – buildings buffered + setback) on non-residential parcels with 5–60 % building coverage. The Layer B output is by construction an *upper bound* on parking — it includes drives, side yards, alleys. NAIP-segmentation refinement (Phase 4 Layer C in the project plan) is deferred. The parking-area totals in this paper should be read as candidate-for-parking, not proven parking.

7. **Single-year snapshot.** We use TY2025 appraisal data and the FY2026 tax rate. The City updates its tax rate annually and HCAD reappraises every year. Reruns are straightforward (all pipeline scripts are idempotent and parameterized on tax year), but this paper is a TY2025 snapshot.

8. Reproducibility

The analysis is a numbered Python pipeline plus a static frontend. All intermediate artifacts are checked into the repository's `data/` directory structure (gitignored for size; reproducible from public sources).

```
pipeline/
  01_fetch_hcad.py           Download HCAD annual public-data files
  02_fetch_osm.py            Pull amenity=parking polygons via Overpass
  02b_fetch_buildings.py     Microsoft USBuildingFootprints, bbox-filter to Houston
  03_fetch_overlays.py       COH boundary, super-neighborhoods, TIRZ, METRO, FEMA…
  04_build_base.py           Parcel base layer, COH centroid filter, HCAD attr join
  05_classify.py             Vacant / parking-primary / low-improvement bucket tag
  06_detect_parking.py       OSM clip + footprint subtraction → parking polygons
  10_tax_compare.py          Current-uncapped tax + revenue-neutral LVT solver
  11_publish_tiles.py        Tippecanoe → PMTiles for the static frontend
```

Each script logs structured JSON-per-line records to stderr at every key boundary (download size, sha256, row counts, validation gates), so a rerun's output can be diffed against this paper's numbers without ambiguity. The solved LVT rate, certified-levy ratio, and per-bucket breakdowns in §6 are read directly from `data/final/summary_aggregate.json`, which `10_tax_compare.py` emits.

The interactive map at https://powerfundamental.com shows the same per-parcel results visually. PMTiles are served from Cloudflare R2 (custom domain `powerfundamental.com`); the frontend is hosted on Cloudflare Pages.

Repository: source code accompanying this paper, available on request. License: source MIT, data inherits the upstream license of each input source (HCAD public-data terms, OSM ODbL, Microsoft Open Data Commons-By, FEMA public domain).

Working draft, 2026-04-26.