GUBMENT POLICY WHITEPAPER NO. 2 · SERIES GBMT-2 · FILED 2026-08-03

Housing: why doesn't legalizing it build it?

A feasibility assessment of the leading "fix housing" theories — built around a design detail that separates reforms which produced tens of thousands of homes from reforms that produced dozens, a shortage number that turns out to be an assumption wearing a measurement's clothes, and 278,345 public-housing units that are already eligible for federal capital and operating funds.

READ 90 SECONDS 10 MINUTES FULL RECORD

CONTENTS / RECORD +

ABSTRACT

This inquiry set out to name the binding constraint behind America's housing crisis and score the leading fixes against it. Three findings organize everything else. First, the national "housing shortage" number is not a measurement: each publisher’s residual — their entire pent-up-household apparatus — is smaller than the swing from one percentage point of their own vacancy assumption — one percentage point of that assumption is worth about 1.59 million units, and the published estimates' entire spread fits inside it. Second, legalizing a use is not legalizing a building: reforms that added real buildable envelope (California's prescriptive ADU laws, Oregon's middle-housing law, Houston's lot-size cut) produced tens of thousands of homes, while reforms that raised the permitted unit count but left the envelope and local discretion intact (California's SB 9, Austin's HOME) produced almost none — an order-of-magnitude difference on administrative data, with the mechanism read out of each city’s own ordinance record. Third, the constraint map keeps pointing away from where the public argument looks: homeowner opposition is real but not about home values; the swing vote that decides state bills is construction labor; insurance is quietly eating the operating side of every affordability program; and the "cap" on public housing has 278,345 units of headroom this project verified from HUD's own file. A nine-architecture scorecard, red-teamed and then independently re-scored (three cells corrected, then thirteen more, the draft headline withdrawn twice), lands on one evidenced leader rather than a ranking — and on two constraints no proposed instrument addresses at all. A Phase 2 steelman then found the leader's lead is one cell wide, that the board's missing *rows* matter more than its cells, and that the two randomized trials of the instrument this filing cautioned against are absent from its record.

PART 1

The number that isn't: the housing shortage is an assumption in a trench coat

KILL CONDITION FIRED

1.0–8.2M
PUBLISHED SHORTAGE ESTIMATES

1.59M
UNITS PER 1PP OF ASSUMED VACANCY

3%
RESIDUAL REPRODUCING FREDDIE MAC'S 3.7M

66.5%
OF VACANT UNITS ARE OFF-MARKET

Nearly every mainstream shortage estimate belongs to one method family: pick a "target" vacancy rate, subtract reality, call the difference the shortage. Running that model on committed Census data, one percentage point of assumed target vacancy is worth about 1.59 million units — and feeding each publisher's own stated target into the bare model reproduces Freddie Mac's famous 3.7 million to within 116,000 units from the assumption alone. The published range (5% to 12% targets) spans more units than the entire spread of published estimates. Two federal surveys agree on how many housing units exist to within 0.06% and disagree on how many are empty by over a million.

shortage_sensitivity.py, committed pipeline · CRS IN12628 · ACS 2024 / CPS-HVS direct tabulation

> PLAIN TALK
>
> Nobody measured the housing shortage. Every big number you've heard is mostly a choice about how many empty homes America "should" have, dressed up as a count. That doesn't mean there's no shortage — it means the number is an opinion, and the opinions disagree by millions.

PART 2

Legal on paper, unbuildable in fact: the design detail that decides it

25,000+
CA ADU PERMITS/YR BY 2022

38
SB 9 UNITS APPROVED, LOS ANGELES

51 → 8 → 18
AUSTIN 3–4-UNIT-CLASS UNITS, 2023 → 2024 → 2025

Same state, same era, two designs. California's ADU laws — prescriptive, ministerial, envelope-adding — took production from about 540 permits in 2016 to over 25,000 a year, a fifth of the state's housing output. SB 9, which legalized more units per lot but left local standards and discretion intact, produced 38 approved units in Los Angeles and zero lot splits; eight of thirteen surveyed cities approved none. Austin repeated the experiment independently: its HOME reform legalized three units per lot while leaving the buildable envelope at 0.4 FAR, and the three-to-four-unit permit class *declined* after the reform legalizing it — 51 units in 2023, 8 in 2024, 18 in 2025, still roughly a third of the pre-reform year two years on. Querying the same database by unit count, the HOME three-unit path produced about 15 permits in two years, against 233 permits at two units. Both come from the city's own permit database, not press coverage. Two limits travel with them: that Austin books a HOME three-plex into the three-and-four-family class is confirmed by that unit-count query, not against Austin’s written permit-classification guidance; and the same series carries a classification break — the two-family class rises 156 → 532 while the secondary-apartment class collapses 248 → 10 over the same window — so the small-scale totals cannot be read as pure new production.

Terner Center ordinance review · City of Austin open data (austin_decompose.py, committed) · phase0-findings §3

Every reform that added buildable envelope produced homes. Every reform that raised the unit count while holding the envelope produced almost none — two cities, two states, administrative data, mechanism documented in each city's own ordinance record.

PLAIN TALK

Cities have learned to say yes in a way that still means no: legalize triplexes, keep the height and size limits that make triplexes impossible to build, take credit either way. The reforms that actually work are boringly specific — fixed standards, no discretionary review, real buildable room.

PART 3

Who actually blocks housing — and it isn't about their home values

MECHANISM REFUTED — AT MODERATE CONFIDENCE

The standard theory says homeowners block housing to defend their asset values. The evidence runs the other way: upzoning tends to *raise* the value of the upzoned lots, and homeowners oppose it anyway — which means the binding motive is something else (neighborhood character, parking, control), and instruments built to compensate asset losses are aimed at a mechanism that isn't operating. What actually killed the recent wave of state bills wasn't the tenant-vs-developer fight the press covers: New York's Housing Compact and Colorado's 2023 preemption bill both died to suburban home-rule revolts inside the governing party. And the swing vote with the best-documented arc is construction labor — California's Building Trades killed the 2018 and 2020 upzoning bills, then endorsed or went neutral on the 2023 and 2025 ones after enforceable labor standards were written in. The bills passed with the tenant coalition still opposed.

`ws07-findings (H7.1) · ws08-findings: SB 827/50/423/79 arc, NY Compact, CO SB23-213 · four documented state formulas`

PLAIN TALK

The people who show up to fight new housing aren't protecting their property values — the data says new zoning usually raises them. And the fights that decide state housing bills aren't the ones on cable news: they're about union labor standards and suburban towns' right to say no.

PART 4

The workers nobody's plan supplies

2.6–63×
REQUIRED WORKFORCE GROWTH VS. TREND

26.3%
CONSTRUCTION WORKFORCE FOREIGN-BORN, 2024

15–23%
UNDOCUMENTED SHARE (MODELED BAND)

< private avg
CONSTRUCTION WAGE GROWTH 2020–26

Closing any defensible shortage band requires residential-construction workforce growth between 2.6 and 63 times the sector's recent trend, depending on the band and the ramp — 3.8 to 31 times on a ten-year ramp. The channel that has quietly absorbed the gap is immigration — a record 26.3% of the construction workforce is foreign-born, roughly one in three in the trades, with the undocumented share modeled at 15–23% — a band assembled from three organizations estimating three different populations (all construction workers, laborers only, Texas only), not three measurements of one quantity. That channel is exposed to near-term enforcement policy, not slow demographics. Meanwhile the textbook signature of a binding labor shortage — a construction-specific wage premium — is absent: construction wages grew fractionally *slower* than the private-sector average over 2020–2026. Hard costs, it turns out, are not the lever in constrained metros at all; land and process are.

construction_labor_model.py, committed · BLS CES direct pulls · NAHB/ACS foreign-born series · Pew/CAP/AIC convergent band

> PLAIN TALK
>
> Every housing plan assumes the workers exist to build it. The math says we'd need construction hiring to run several times its normal pace — and the workforce that would do it is the one most exposed to immigration enforcement right now.

PART 5

The leaky plumbing: vouchers landlords refuse, a subsidy nobody can price

67–78%
VOUCHER DENIAL, UNPROTECTED MARKETS

3 in 4
ELIGIBLE HOUSEHOLDS RECEIVING NOTHING

~83¢
LIHTC EQUITY PER CREDIT DOLLAR

none
FEDERAL AGENCY ASSIGNED TO LIHTC COST DATA

Federal rental assistance is rationed, not an entitlement: three in four income-eligible households receive nothing. The flagship demand-side instrument leaks where it matters most — landlords refuse vouchers at 76–78% rates in the two study markets with no source-of-income protection, and at 67% in a third with partial protection, and refusal is *highest in low-poverty neighborhoods*, exactly where the voucher's mobility purpose points. That leak is fixable, and this filing missed the evidence: a randomized trial in Seattle and King County paired vouchers with search help and direct landlord engagement and moved families into high-opportunity neighborhoods at 53% against 15% for the control group, at $2,670 a household — while a straight increase in what the voucher pays in expensive neighborhoods moved almost nobody. On the supply side, the flagship program is a measurement black box: no federal agency is assigned to maintain and analyze LIHTC cost data, the widely-assumed cost penalty versus market construction has never been tested nationally (and state studies point the other way), and the roughly 17-cent gap between a credit dollar and the equity it buys cannot be decomposed into time value, investor return, and fees from any public source.

CBPP · GAO-18-637 / GAO-24-107064 · Novogradac equity pricing · ws05-findings · Bergman/Chetty/DeLuca/Hendren/Katz/Palmer, *AER* 114(5) 2024 (added by the Phase 2 steelman)

> PLAIN TALK
>
> The biggest rental-help program turns away three of four eligible families — but a cheap package of search help and landlord outreach turns that around in a controlled trial, and this filing didn't find that until after it published. The biggest building program can't tell you what its units cost — because Congress has never assigned anyone in the federal government to maintain and analyze the figure.

PART 6

Insurance eats the subsidy

2×

MULTIFAMILY INSURANCE PREMIUMS, 2021→24

~45%

PREMIUM INCREASE, 2023→24 ALONE

0

PROPOSED INSTRUMENTS THAT ADDRESS IT

Multifamily property insurance premiums roughly doubled between 2021 and 2024 — more than six times the rise in CPI over the same period, and up ~45% in 2023–24 alone. The asymmetry is the finding: a market-rate owner passes an insurance shock through to rents; a rent-restricted owner cannot — rent restriction is the program's core mechanism — so the same shock that market housing passes along, affordable housing absorbs out of net operating income until maintenance defers or the asset is lost. Every architecture in the standard reform debate is a capital-side or demand-side instrument; none touches the operating side. That diagnosis comes straight from the workstream evidence. The candidate fixes (insurance pooling, backstops, parametric reinsurance) do not — they are this desk's suggestions, labeled as such.

PLAIN TALK

While everyone argues about building new affordable housing, the affordable housing we already have is being squeezed by insurance bills it legally can't pass on. None of the nine reform architectures this filing scored addresses it; whether any bill in Congress does is a question no workstream searched.

PART 7

The cap that isn't: 278,345 public-housing units already qualify for federal funding

CLOSED AT PRIMARY SOURCE

278,345

UNITS OF FAIRCLOTH HEADROOM, HUD'S OWN FILE

~25%

HELD BY JUST FIVE PHAS

1.4M → 878k

PUBLIC HOUSING PEAK → TODAY

The Faircloth limit is usually described as a ban on new public housing. It is actually a condition on two federal funding streams: a housing authority may not spend Capital or Operating Fund money building above the unit count it held on October 1, 1999 — and the statute's own subparagraph (B) lets it build above that line anyway, just without extra formula money. Because the country demolished so much public housing since (HOPE VI alone removed 93,295 units), that funding line now sits 278,345 units above the standing stock. This project pulled HUD's own published Faircloth List and summed it directly — twice, independently, with a committed pipeline whose identity check caught a 13-unit inconsistency inside HUD's own file. A quarter of the headroom sits in just five housing authorities. The binding constraint on public housing is money and capacity, not legal authority — and the fall from a 1.4-million-unit peak to 878,000 units standing today is a completed durability test any revival plan has to answer.

42 U.S.C. §1437g(g)(3) · HUD PIH Faircloth List (04-03-25), pull_faircloth.py, committed · CRS RL32236 (HOPE VI) · ws09-findings

PLAIN TALK

Building more public housing already qualifies for federal funding — about 278,000 units' worth, verified from HUD's own spreadsheet. What's missing isn't permission. It's money, and a plan for surviving the next Congress that hates it.

PART 8

What other places prove — and what they can't promise

TRANSFERABILITY, HONESTLY

Tokyo out-built all of California — 124,810 starts against 111,760 permits in comparable years, roughly 3.2× per capita, and the standard caveat about disposable Japanese houses didn't survive root-tracing (the two housing stocks turn over at comparable rates, making Tokyo's advantage larger, not smaller). Auckland's celebrated upzoning result is contested — every favorable estimate runs through one author, with three confounders inside the window. Vienna's social housing is real and enviable, but its enabling condition — a land bank acquired over a century starting when rent control had crushed land values — is non-portable, and Vienna's own peer-reviewed literature shows land scarcity now constraining the model at home. Singapore's system rests on compulsory land acquisition at below-market prices — and it is the below-market pricing, not the compulsion, that the Fifth Amendment's just-compensation requirement forecloses. Sweden's queue is the cautionary tale done precisely: nine years average in Stockholm, 15–25 in prime districts, with a documented black market. And the American postwar miracle was a peak-year artifact — 1950's ~1.95 million starts, not a decade-long average, on financing built atop explicitly discriminatory underwriting no one should copy.

TMG / Census BPS direct pulls · anchor 6 independence audit · Kadi et al., Housing Policy Debate · ws09-findings

> PLAIN TALK
>
> Every country everyone cites is real — and every one of them got there with something America doesn't have: Tokyo's national zoning, Vienna's century-old land bank, Singapore's power to take land cheap. The lesson isn't "copy them." It's which specific pieces travel.

PART 9

The scorecard, twice-checked and then steelmanned: the lead is one cell wide

ARCHITECTURE	SUPPLY	RELIEF	MOBILITY	STABILITY
ENVELOPE-PRESCRIPTIVE PREEMPTION	4	3	3	3
FEDERAL MONEY-FOR-UPZONING	2	3	3	3
LIHTC EXPANSION	3	3	3	2
VOUCHER ENTITLEMENT (BARE)	1	3	1	3
FAIRCLOTH SOCIAL HOUSING	2	3	3	3
LAND-VALUE TAX	3	3	3	3
INDUSTRIALIZED CONSTRUCTION	1	2	3	3
CASH (COMPARATOR)	1	3	3	3
HOUSING-FIRST	—	—	—	—

Nine seeded architectures, four objectives, six weightings — and two full adversarial checks. A red team caught the first draft scoring its favorites as idealized composites; then a second scorer, blinded by instruction rather than by structure, caught the correction itself still enforcing the no-evidence-no-score rule in only one direction — nine cells sat below neutral on mechanism hunches, three cells missed evidence that was sitting in the record. Reconciled, the honest result: this record produced above-neutral evidence for exactly one instrument — envelope-prescriptive legal reform — cautions industrialized construction, whose scaling thesis the record refutes, carries one genuinely mixed verdict (Faircloth social housing: real legal authority, adverse operating history), and is honestly agnostic on the rest. Most of what looked like ranking was evidence-density. Housing-First stays unscored. Both scorers, working blind of each other, independently agreed the record then held nothing to score it on; a late evidence pass (§12, 2026-08-06) has since supplied one, and the row moves only after a fresh matrix pass.

A third check — the Phase 2 steelman, 2026-08-10 — found that "leads every weighting" is thinner than it reads, and withdrew one caution. Twenty-three of this board's thirty-two cells are neutral, and envelope-prescriptive preemption holds the only cell above neutral anywhere on it — so its lead follows from the shape of the matrix rather than from running the weightings, and it is worth a quarter of one point. That cell is also the one our two adversarial checks split on: the red team argued it down to 3, the blind second scorer marked it 5, and it was reconciled to 4. Read at the red team's 3, envelope-prescriptive preemption leads none of the 1,771 weightings we swept and ties the land-value tax, the row this filing itself calls an ignorance artifact. The arithmetic is now committed and reproducible. And the caution against the bare voucher does not survive: HUD's own randomized experiment across 2,282 families in 12 cities cut homelessness from 38% to 17% using an ordinary housing voucher, and a second trial fixed the voucher's mobility failure for $2,670 a household. Neither study appears anywhere in this filing's record. The caution is narrowed to what it was always about — a voucher builds no housing — and withdrawn as a verdict on the instrument.

`9 architectures × 4 anchored objectives × 6 weightings · red team + independent blind re-score, both reconciled and committed · Phase 2 steelman with the ranking swept over every weighting (ws10/rank.py, committed)`

STEELMANNED; ONE CAUTION WITHDRAWN

APPENDIX

The honesty box

The shortage number is not derivable, and we say so instead of picking one. The filing's own kill condition fired: the headline statistic of the entire debate is parameter-driven, and our contribution is the sensitivity surface, not a better point estimate. **The homelessness track ran late and thin.** The protocol declared it parallel and it stayed parallel through publication; a bounded evidence pass finally ran on 2026-08-06 (§12), finding Housing First supported on housing stability for a narrow high-needs population and *refuted* on paying for itself. Housing-First is still unscored rather than invented — the record's largest acknowledged gap. **The scorecard's first draft flattered its favorites.** A red-team pass caught idealized-composite scoring for the two leaders, an uncited cell above the evidence floor, and an untested neutral weighting that ties the top three; all corrections are in the record, with the sensitivity reading published. **The mobility objective is nearly unevidenced across the board** — the only direct mobility evidence we found was the voucher's, and it is negative; the Phase 2 steelman then found the randomized evidence that it is also reversible, which we had not searched for. **We scored the instruments we thought of on day one.** The nine-architecture list was written before the first workstream ran and no workstream ever added to it — so the board carries no tenant-protection or anti-displacement instrument even though one of its four objectives *is* stability, and no preservation instrument even though Part 6 is about a constraint that destroys existing affordable units. Our red team, our second scorer and our Phase 1 fact-check all audited the cells; none audited the rows. **We cautioned against the voucher without reading its randomized evaluations.** HUD's Family Options Study and the CMTO trial appear nowhere in this record; both were found by the Phase 2 steelman, and both cut against what we published. **The ownership-vs-rental subsidy claim came back indeterminate** on strict two-source grounds, supported only on the narrow reading, by a margin too large to be rounding. **One analyst, one desk.** Criteria were pre-registered and every verdict cites its workstream, and the second scorer has now run (2026-08-06) and moved thirteen cells. Its blinding was instruction-level, not structural — the weaker of the two standards this project uses. When elder care was re-scored under structural blinding on 2026-08-10, 25 of 40 ranked cells moved and a published headline was withdrawn. There is no reason to assume this board would hold up better.

THE RECEIPTS · 13-workstream protocol · Phase 0 gate with a GO verdict and a fired kill condition · pre-registered anchor table preserving broken priors (anchors 1, 6, 10; anchor 9 closed at primary source) · deviations log (76 entries) · committed data pipelines (ACS, CPS-HVS, BPS, BLS, Austin admin data, HUD Faircloth) · adversarial red-team pass (three cells corrected) plus an instruction-blinded re-score (thirteen more) · a Phase 1 verification pass against primary sources and a Phase 2 steelman that withdrew one caution · ranking arithmetic swept over every weighting, committed · all public.

CITE · Gubment. "Housing: Why Doesn't Legalizing It Build It?" Policy Whitepaper No. 2, Aug 2026. gubment.com/housing · Sources & data: gubment.com/housing/sources