Frozen Frontage

The Comp Trap: How Building Finance Keeps Rents High and Offices Empty

Working paper — draft

June 2026

Abstract

Downtown Seattle is a third empty, and the rent has not moved. In a working market someone undercuts the building next door and fills the floor, but that is not happening here. This paper explains the financial trap that prevents it, where filling your own building drives down its own value first, and every other building's with it. A permanent, demand-driven revaluation (Gupta, Mittal, and Van Nieuwerburgh 2026) explains why downtown is empty; it does not explain why owners keep asking rents no one will pay instead of cutting to fill the space and earn the revenue.

That second fact is an equilibrium, and the reason is a single cheap lease: sign one to fill a floor and you print a low comparable that the income-approach appraisal uses to mark down your own building first and hardest (and every neighbor's a little), and it is your own fresh comp that breaches your covenant first, so no one signs it. The result is a cartel-like freeze: it rests on no agreement and collects no rent, because the space sits empty, yet it holds asking rents high. That makes it a coordination failure for government to fix, not an antitrust case. We call it the **comp trap**. The held face rent does more than guard the owner's collateral: because the personal guarantee a small tenant must give is sized to the *face* they sign, an inflated face raises the cost of *taking* the space, so the floors stay dark. Demand owns the bulk of the vacancy; the comp trap holds asking rents far above it, and plausibly keeps some extra space empty too—most of all at the most heavily indebted buildings. The 2025 record is the mechanism in motion—the most-levered owners conceding first, building by building, and leasing rebounding where the low comp has already printed (San Francisco)—so the break is the cure when the building changes hands, and policy's job is to force that and stagger it.

The fix is pro-market: a *flat* vacancy tax only re-times default, but a credibly *escalating* idle-time tax forces a staggered, building-by-building exit, with prudent-use assessment as the base-stabilizing complement that protects homeowners. The falsifiable test keys not on the vacancy level, which demand dominates, but on one specific comparison: when a neighbor falls into distress, the owners who hold their asking rents highest in response are the most tightly-covenanted ones, even after controlling for their own leverage (§8).

1 Introduction

A real demand shock explains why downtown Seattle is empty; this paper explains the separate puzzle of why owners hold asking rents high instead of cutting to fill—and why the fix has to be a vacancy tax that escalates the longer a building sits empty, not a flat one.

Roughly one in three office floors in downtown Seattle sits empty, and the rents landlords post for them have barely moved. In a working market an empty floor is an opportunity: someone undercuts the building next door, signs a tenant, and the space gets used. That is not happening here. Filling a building has become the move that hurts its owner, because the cheap lease it takes to fill it marks down its own appraised value first and hardest (the bank's collateral), and every neighbor's a little, so everyone holds instead. That backwards incentive is what this paper explains, not the emptiness itself.

The one-line version: cut your rent to fill a floor and you create a cheap *comparable lease* that the standard appraisal uses to mark down your own building first and hardest, and every neighbor's a little, tripping your own covenant before anyone else's—so no one cuts, and the block stays frozen at high asking rents even as it empties. We call it the **comp trap**. Structurally it is a bank run in reverse: in a bank run the disaster is everyone rushing the exit at once; here it is everyone refusing to move, each owner safe only as long as the others sit still too.

Much of the emptiness is well understood. Remote work permanently reduced office demand, and the post-COVID downturn is a structural revaluation, not a cycle (Gupta, Mittal, and Van Nieuwerburgh 2026; §2.1): the twenty tallest towers in Seattle's core have lost about $5.6 billion in assessed value since 2021, and Blackstone agreed to sell the US Bank Centre for 54% less than it paid in 2019. We take that demand collapse as the baseline. What it does *not* explain is the rest: posted asking rents that have barely moved even as effective rents and values crater, and owners holding space dark for years rather than cutting to clear.

Why would an owner do that? Because a building is valued by its income—roughly, its rent divided by a market *cap rate*—so its value moves with the rent on recent comparable leases nearby. The discounts landlords actually give are buried in side deals (free months, build-out budgets), so it is the high *headline* rent that lands in the comp databases and contaminates the block. Signing one cheap lease therefore marks the owner's own building down first and hardest, trips its loan-to-value covenant at the next refinancing or appraisal, and re-rates every neighbor the same appraisers value next. Holding is each owner's best response—and a best response to everyone else's holding, since the option to lease later is worth more while the comp set stays clean.

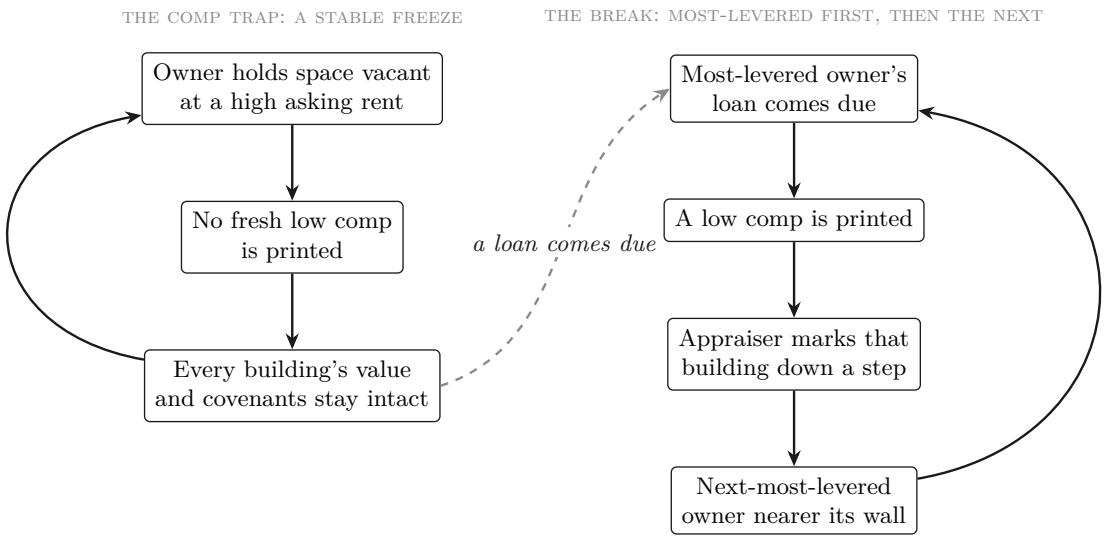

Figure 1: The comp trap. While every owner holds, the high comp set protects everyone's collateral and the block stays frozen (left). One forced sale prints a low comp and reprices that building; the next-most-levered owner is forced out when its own loan matures, and the freeze gives way building by building (right)—the pattern observed in downtown Seattle, 2024–2026.

While everyone holds, no cheap comp prints, no value is marked, and no covenant breaks: the freeze sustains itself. Every cartel you know of rigs prices to collect *more*; this one holds asking rents high and collects *nothing*, because the space is empty. It rests on no agreement and prosecutes no one—just every owner best-responding to the same loan covenant—so it is a coordination failure for tax policy to fix, not an antitrust case (§3.2). What finally breaks it is not nerve but the calendar, and downtown Seattle is watching it happen: Martin Selig, long the largest owner downtown, ran out of runway first and defaulted building by building, one tower at a time as each loan came due—the most-levered owner forced out first, by default, not a discount, exactly as the model predicts (§7). The policy follows, and is pro-market throughout

(§6): a flat vacancy tax only re-times that default, but an idle-time tax that *escalates* on each building's own clock forces a staggered exit and restores the price discovery the freeze suppresses.

Two further pieces complete the picture: the held face rent also raises a tenant's *own* cost of taking the space—a personal guarantee sized to the face—so the trap reaches the vacancy level at the margin, not just the wedge (§3.3); and the 2025 record shows the break already underway, leasing rebounding where the comp has printed, which §7 develops.

Contributions. (1) An appraisal-comp microfoundation for a cross-property collateral externality. (2) Asking-rent setting as a coordination game in which *held rents amid vacancy* are the bad equilibrium. (3) A demand-side channel—held *face* rents, to which a small tenant's personal guarantee is sized, raise the cost of *taking* the space—that ties the trap to the vacancy level at the margin, not just the wedge. (4) A Seattle calibration on public CMBS and assessor data. (5) The leverage-dependent default boundary and the vacancy-tax trap. (6) A policy playbook and a falsifiable empirical test (§8).

2 The Seattle Setting

The demand collapse is real and structural—but the comp trap's fingerprint is the separate fact that face rents are held high while values crater, visible in the public record alongside the leverage anchors and the staggered, leverage-sorted defaults the model predicts.

2.1 The demand shock is real and structural—concede it

Downtown vacancy has spiked twice this century, and the two ended very differently. The 2008–09 financial crisis pushed downtown vacancy to ∼18.9% and it then *recovered* to ∼7% by 2017–2019; COVID pushed it from that ∼7% to 36.5% by Q1 2026 and it has *not* recovered. One spike came back; this one keeps rising—which marks the episode as structural, not cyclical.

This is not a normal cycle. Gupta, Mittal, and Van Nieuwerburgh (*AER* 2026) estimate a long-run office-value decline approaching 46% in New York and ∼$557B nationally, driven by permanently lower expected cash flows, with lower-quality buildings becoming "stranded assets"—explicitly long-run, not cyclical. That rival is our baseline: it owns the bulk of the vacancy level, the value collapse, and the failure to recover. The comp trap earns its place on the residual the demand story cannot produce—the held rents of §2.3, and the leverage-sorted increment of vacancy they keep frozen in place.

Why the shock cut so deep. Remote work is the trigger; the *severity* reflects what kind of downtown it hit—a near mono-use office core with little residential base and few non-office tenants. The downtowns that recovered worst are precisely these office monocultures: a cross-city comparison ranks Seattle's and San Francisco's recoveries near the bottom, the slowest belonging to information- and finance-heavy cores (Forouhar et al. 2025). Mixed and residential districts held up—San Francisco's residential neighborhoods refilled even as its Financial District did not. An office-only block has nothing to recover to after 7pm and no other tenant to reprice into. This does not compete with the comp trap, which explains the held-rent *wedge*; it explains why the shock bit so hard, and why *conversion*—not just repricing—is the durable fix (§6): every floor turned to housing rebuilds a resident base mono-use zoning never allowed, so people can work remotely *from* downtown rather than abandon it.

2.2 Among the worst in the nation

On a downtown/CBD basis Seattle (∼36%) is in the worst tier, just below Denver and roughly double Manhattan; on the broader *market-wide* basis (∼33%) Cushman & Wakefield ranks Seattle the single most vacant major market in the country (Figure 2). The two figures cover different geographies.

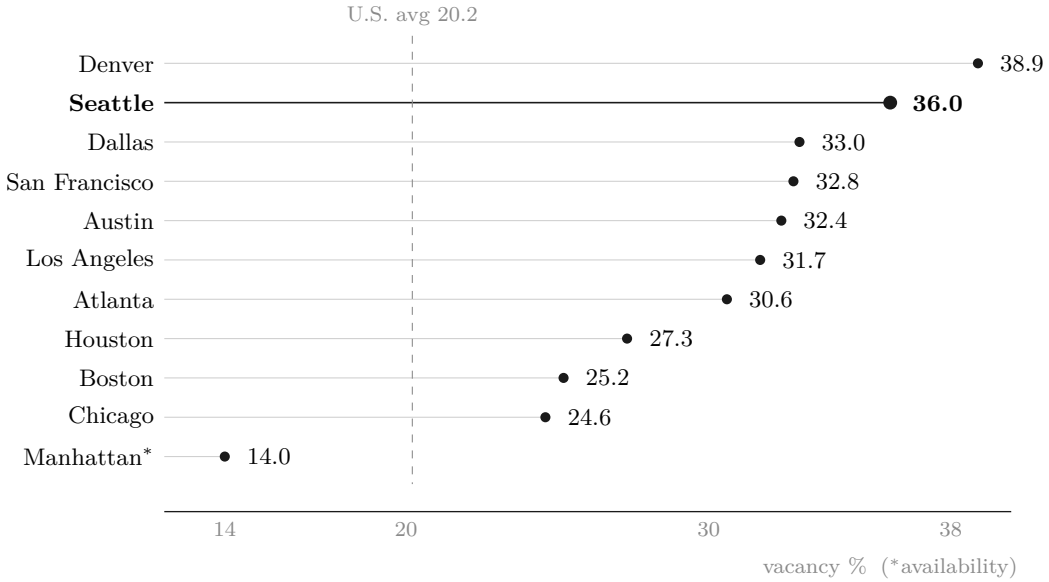

Figure 2: Downtown/CBD office vacancy, major U.S. cities (Q4 2025–Q1 2026). Seattle sits in the worst tier, $\sim 2\times$ Manhattan. Sources: CBRE, C&W, Colliers, Partners (per city).

2.3 The discriminating fact: sticky face rents while values collapse

This is the comp trap's signature, and the only fact that separates it from the structural-demand baseline. A permanent demand shock clears by lowering *posted* rents; the comp trap holds them. Seattle "leads the country in falling office rents," but the decline shows up far more in *effective* rents (after concessions) than in *face* rents—TI allowances run up to \$100/sf and free rent near nine months, even as top-tier face rents hold. This is the wedge the mechanism predicts. Appraisers *do* net concessions when they can see them—but concession packages are private and reported late, so the figure that propagates across the block, through the comp databases a neighbor's appraiser pulls, is the held *face* rent. A landlord posts \$50/sf, then privately hands back 25–40% of the economic rent through long free-rent periods and oversized tenant-improvement budgets (running roughly a third above pre-pandemic levels, per CBRE). Their effective rent falls, but the contaminating headline comp never hits the record. Industry participants describe the motive openly: structure the concessions so the headline rent stays up, and the *cross-building* comp every other appraiser pulls—and the next *external* refinancing or sale appraisal—never marks the collateral down. (The owner's own current lender already holds the abated lease and sees the true in-place economics—though its DSCR and minimum-rent covenants are typically written on *base* rent, so the held face props the owner's own loan too; what the face *additionally* protects is the value every other valuation is built on.) The high face rent is often not purely the owner's choice: a lender's major-lease-approval right can block a specific below-threshold signed lease, so the held rent is partly the *lender's*—a covenant that relocates the externality rather than removing it (§3.2). Each owner does this individually, no coordination required, but the aggregate is a face-rent signal held aloft while the rent actually paid falls—the price-discovery failure the comp trap describes. Appraisal smoothing (Quan and Quigley 1991) makes the eventual re-rate partial and lagged rather than instantaneous, which lengthens the freeze rather than dissolving it.

When the re-rate finally comes, it appears in King County Assessor income-approach valuations, which we read at the parcel level (Table 1).

Table 1: The comp re-rate in the public record (King County Assessor, eRealProperty).

Building	2021 assessed	2026 assessed	Decline
US Bank Centre (1420 5th Ave)	\$573.7M	\$251.2M	−56%
1201 Third Ave (former WaMu Tower)	\$699.7M	\$324.3M	−54%
(of which one single-year cut, 2024→2025)			−37%
Amazon Doppler Tower	—	—	−62%

2.4 The staggered default the model predicts

Martin Selig, long downtown's largest owner, ran up >\$800M in default across 18 buildings: the Federal Reserve Building (∼82% vacant) and 400 Westlake (∼94% vacant) handed to Acore on a \$240M default; seven buildings backing \$239M of CMBS transferred to CW Capital for \$120M (∼50% recovery). This is the *break* path of Figure 1 realized.

2.5 The cross-section: above-market by design

The held-rent wedge is not uniform across the block, and its variation is the trap's most visible fingerprint. The mechanism predicts *sorting*: owners who can price at their building's market-clearing rent fill, while owners who must hold asking *above* clearing—to avoid printing the low comp that breaches their covenants and marks the comp set—sit empty. Vacancy should rise in the asking premium over clearing, splitting the cross-section into a market-rate cluster (low premium, low vacancy) and a comp-trapped cluster (high premium, high vacancy); the slope between them is the wedge made visible (Figure 3). The tier data take this shape. In Seattle, trophy towers ran ∼10% vacant in Q4 2025 while broad Class A ran ∼32% (Cushman & Wakefield); nationally CBRE's prime (∼14%) and non-prime (∼19%) series split the same way. The divergence runs through *posted* rents as the trap predicts: across ∼4,350 new-lease comparables in twelve markets, CBRE finds top-tier face rents holding or rising since 2023 while lower-tier face rents fell only slightly even as their effective rents fell more—face held aloft in the weak tier while concessions widen, the held wedge in aggregate. Bellevue's Bravern Office Commons went dark after Microsoft's exit and still carries full asking rent, a current owner holding the line; Selig's Federal Reserve Building (∼82% vacant) and 400 Westlake (∼94%) carried full asking rents into near-total emptiness before he defaulted onto their lenders, who hold the price still.

A clarification the model demands: the relevant premium is *not* the trophy-over-commodity quality premium—prime space genuinely commands more—but asking rent over each building's *own* market-clearing rent, holding quality fixed. The quality premium is a level effect the building and submarket fixed effects of §8 absorb; the own-clearing premium is the wedge the cross-partial isolates. The split is the one §3 derives: owners who can price at their own clearing rent fill, while owners who must hold above it—to avoid printing the comp that breaches their covenants—sit empty, and the gap is widest where leverage and comp-set distress are both high (the §8 prediction). The lease-up hazard falls in the premium, which the §8 hazard model measures directly.

2.6 The homeowner burden shift—mechanism and honest limits

Washington runs a *budget-based* property tax: under the ∼1% levy-growth cap (RCW 84.55), each district fixes a dollar levy first and the assessor back-solves the rate = levy ÷ total assessed value. Assessed value sets only each property's *share*. So when the commercial base shrinks, the same levy redistributes onto the rest and the residential share rises, the classic property-tax incidence channel (Oates 1969; Table 2).

Figure 3: The cross-section of the wedge. Each point is a building or tier: vertical position is vacancy, horizontal is how far asking rent sits above that building's *own* market-clearing (effective) rent. Market-rate owners price at clearing and fill (bottom-left); above clearing, buildings sit empty (top-right), whether a current owner holds the line to protect a live covenant or a defaulted building sits un-repriced under its lender. The vacancy figures are measured (sources in §2.5). The horizontal premium is modeled: building-level asking-minus-effective rent is proprietary, and measuring it directly is exactly what the §8 test does. So the claim is the upward slope, not the exact coordinates.

Table 2: The budget-based mechanism, stylized. The city collects $100 both years; the homeowner's bill rises 33% with no change in the home's value.

	Commercial	Residential	Total	Rate	Homeowner share
Year 1	$50	$50	$100	1.00	50%
Year 2 (commercial halves)	$25	$50	$75	1.33	**67%**

What this does and does not establish. The mechanism is exact, but the stylized halving overstates Seattle's actual office channel. King County's residential share did climb, from ∼65% to 83%, but that is mostly *home appreciation*, not the office collapse: commercial value *overall* grew over 2017–2024 ($35.2B→$58.9B), even as *major office buildings* shrank from 7.0% to 3.7% of the county base. The claim we *do* make is directional and about policy: under this system, a *flat*, block-wide tax that forces commercial values down all at once (§3.4) redistributes the levy onto homeowners, while an instrument that keeps the base on the rolls (§6) does not.

3 The Mechanism

We build the comp trap from the ground up—why one owner won't cut even in isolation, why that couples the block into a freeze, why the held rent keeps the space empty*, why a flat tax only re-times the default, and why a tax that escalates on each building's own clock is what breaks the freeze. The formal model is in Appendix A.*

3.1 Why one owner won't cut

Take a single owner near a covenant test—an empty floor, a loan against the building—and forget the neighbors for a moment. Cutting the asking rent to fill the floor backfires on this owner directly: the signed lease is the freshest, full-weight comparable against the owner's *own* building (valued by the income approach off recent comps), so it marks the owner's own collateral down and trips the loan-to-value covenant at the next mark. Those marks are *event-driven*, and the

timing is what makes holding safe: a covenant is tested at refinancing, maturity, or a transfer to special servicing, not continuously, so while the owner simply holds, the markdown does not yet exist—paper value stays stale until someone is forced to a marking event and prints the first low comp. So holding is self-protective even with no neighbors and no game: the cost of being the one to print the cheap comp—the *contamination stake*—exceeds the *lease-up gain* from filling the floor now. But on its own that yields only a *level* of holding; the coordination layer of §3.2 is what turns it into a block-wide freeze that unwinds building by building, and what the *escalating* (not flat) tax is built for. (Appendix A.1 gives the formal environment.)

3.2 Why the block freezes

Now put the neighbors back. Because every building on the block is valued off the same comp set, the cheap lease you print does not only mark *you* down—it re-rates every neighbor whose appraiser pulls it. So holding is also a best response to *their* holding: as long as everyone holds, the comp set stays clean, and the option to lease *later* into a clean set is worth more the more neighbors hold. Asking rents move together: each owner holds because the others do, and the block can settle into either of two outcomes—everyone frozen, or everyone clearing—with a tipping point between them (Appendix A). §6 shows those two regimes empirically, and where the policy moves Seattle's frozen towers (Figure 6).

The freeze is sustained not by an agreement but by the payoff coupling itself. While everyone holds, no distressed comp has printed, so every owner's appraised value—the only value its covenants are tested against—stays high on paper. Whoever cuts first prints the block's sole distressed comp, and because that comp is freshest and full-weight against the cutter's own building while a forced neighbor sale is discounted as a fire sale, the cutter breaches its own covenant before and harder than anyone else. So "all hold" is self-enforcing: a cartel-like freeze that props up *collateral values* instead of rents and collects nothing, because the space is empty. There is no agreement among owners to prosecute: each is simply best-responding to a shared constraint. Fixing it is a job for tax and assessment design, not antitrust.

The lender can be the hub. The lender is itself a coordinating node. When several buildings on a block sit in the *same* CMBS pool or answer to the *same* lender—common where a single originator banked the block—a low comp on any one of them triggers a pool-wide appraisal reduction that subordinates the junior bonds the lender or its first-loss (B-piece) affiliate holds, so its interest is that *nobody* prints it, and its major-lease-approval rights let it enforce that across many owners at once. This is a financing linkage through the lender, not collusion among competitors. It also cuts against the natural guess that a concentrated lender would *internalize* the cross-building externality and tolerate one clearing lease, breaking the freeze: a CMBS pool's special servicer, protecting its junior bonds, does the opposite and *enforces* it, which is why the test below splits the lender effect by type (§8).

The contamination stake—the cost of moving first—is positive precisely when your own fresh lease moves your appraisal more than a neighbor's fire-sale comp would. (The formal decomposition is in Appendix A.2.)

Proposition 1 (The freeze)**.** When the cost of moving first exceeds the rent an owner would capture by filling now—the *no-unilateral-cut condition*, holding for every owner—"all hold high" is a stable (Nash) outcome: no one gains by printing the first distressed comp, because it costs more in covenant damage than the rent it captures and marks the mover down before and harder than its neighbors. The freeze can hold for years, because while everyone holds no comp prints and paper values stay stale; *how* it finally breaks—a determinate, leverage-sorted race—is the subject of §3.4. (Proof: Appendix A.4.)

3.3 Why the held rent keeps the space empty

The story so far explains why owners hold high *asking* rents. It does not yet explain why the space stays *empty*: if the landlord quietly hands back the difference in free rent and build-out budget, the *effective* rent a tenant pays can fall to clearing while the face rent stays high. Why doesn't the floor then fill?

Part of the answer is that the inflated face rent is not a costless fiction for the tenant. A small business usually has to personally guarantee its lease, and the guarantee is written against the *face* rent times the full term—ten months free does not shrink the guarantor's *contracted* exposure—so a high face enlarges the personal liability a business owner must accept to take the space, even when the effective deal is fine. (The tenant's financial statements do *not* carry the face: GAAP rent expense under ASC 842 straight-lines to the effective rate and the lease liability is the present value of actual cash, so a coverage-ratio underwriter sees through it. The binding channel is the contractual guarantee, not the financials—which is also why a "fair-market" renewal, set off prevailing effective comps, leaves the held face as only an opening anchor.)

So the held face rent cuts both ways: the number that shields the landlord is the same number that prices out a credit-constrained tenant. The effect is a *composition* shift rather than a flat addition to vacancy: a concession-loaded deal at a high face still clears for a large, unlevered, no-guarantee tenant, but it prices out exactly the small businesses and new entrants a downtown renewal depends on, so held-face stock self-selects toward emptiness. It fits the split we see—the space that actually leases tends to be priced at a market *face* (§2.5)—and it sharpens the case for the new-tenant carrot of §6. Demand owns the bulk of the vacancy; this is a contributing friction at the margin, not the main driver.

3.4 What breaks the freeze: default, and what a tax does

A levered owner can default rather than lease into a breach. Appendix A.3 sets out the three exits—LEASE, DEFAULT, WAIT—and their payoffs; the one fact that drives everything is that **a vacancy tax falls only on waiting**, never on the two exits (Figure 4). For the deeply underwater owners we feature, leasing and defaulting both wipe out the owner's equity, but leasing into the breach still triggers the covenant penalty that walking away sheds—so the owner defaults, recovers nothing, and the deficiency falls on the lender.

Proposition 2 (Default boundary)**.** There is a market-rent threshold—the *default-trigger rent*—below which default beats leasing, and it *rises with leverage*: the more indebted the owner, the higher the rent at which walking away dominates holding. So as rents fall, the most-levered owners cross into default first. For the deeply underwater, default dominates outright—leasing into the breach triggers the covenant penalty that default sheds. (Closed form and derivation: Appendix A.3–A.4.)

Proposition 3 (Policy trap)**.** Because the tax falls only on waiting—not on leasing or defaulting—an owner already past the default-trigger rent meets a tax steep enough to end the wait by *defaulting sooner*, not by leasing. A flat tax thus only re-times default. The corrective instrument is one that *escalates* on each building's own clock (§6): it staggers the exit—repricing for the lighter-levered, default-and-resale for the over-levered—rather than sparing anyone or synchronizing everyone.

A forced default removes the owner, not always the freeze. A forced default cures the building only when it changes hands to an arm's-length clean-basis buyer who reprices and re-leases it. Passed to the lender or special servicer instead, it can sit frozen exactly as before—Selig's Federal Reserve Building and 400 Westlake, handed to Acore, remain ∼82% and ∼94% empty—which is exactly what the occupancy-keyed tax that runs with the parcel

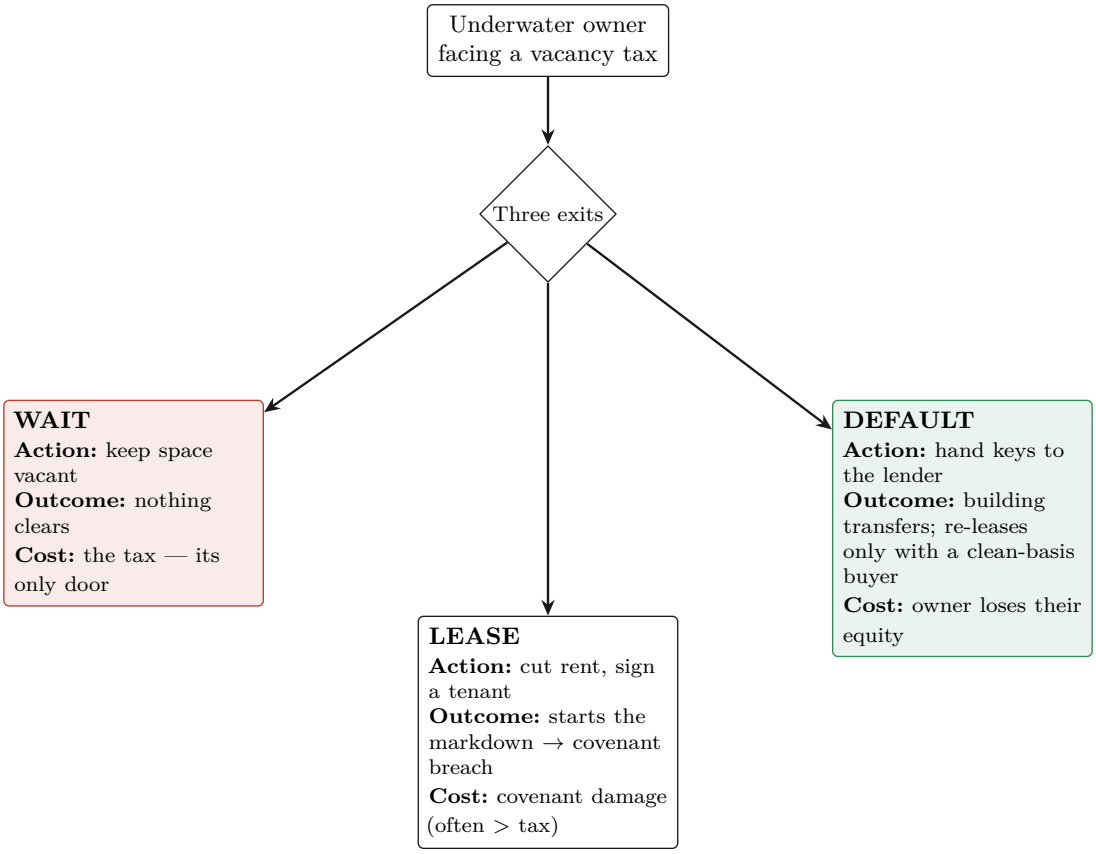

Figure 4: Three exits for an underwater owner, and the tax sits on only one. A *flat* tax big enough to end the **WAIT** swings the owner to **DEFAULT**, not **LEASE**—a foreclosure and a printed-down comp, no lease from the incumbent. Default transfers the building but cures it only with a clean-basis buyer; under a lender it can sit frozen, so the charge accrues with the parcel through default until the building re-leases (§6). §6's escalating, per-building tax forces the choice and staggers the markdowns.

through default is built to stop. None of this is aimed at owners, who were headed for that wall regardless; the tax only sets the date (Prop. 3).

3.5 Welfare: who pays for the dark floors

There are two different harms here, and they call for different responses. The first is the markdown owners inflict on each other: one cheap lease lowers the appraised value every neighbor is borrowed against, tightening their loans. That harm is real, but it is *not* a reason to block the repricing—paper values returning to what buildings are actually worth is the cure, not the cost. What makes a *synchronized* crash worse than a staggered one is that many forced sales at once push prices below even true value (a fire sale; Shleifer and Vishny 1992), so the repricing should happen building by building rather than all at once. The second harm is different in kind, and it falls on people who never signed a loan: the dead frontage. A floor held dark to protect a comp is a floor withheld from everyone who would have used it—the tenant who would pay the clearing rent but is shut out; the small business or new entrant that needs the cut-rate space a downtown depends on to renew itself; and the ground-floor shops, restaurants, and workers who live on the foot traffic the empty floors above them no longer generate (Glaeser et al. 1992). 400 Westlake makes it concrete: about 94% empty at full asking rent, fifteen floors of dark frontage above a sidewalk with no lunch crowd left to feed. How large must the private holding benefit be to rationalize all this? *Option Value and Storefront Vacancy* (Moszkowski and Stackman 2024) works backward from the fact that owners hold rather than lease and recovers a sizable one—which can exceed a single floor's annual rent, since one low comp threatens the entire

building's value. The two harms map onto two instruments: the escalating tax attacks the dead frontage by forcing the space back to use, and prudent-use assessment keeps the base on the rolls so the repricing does not land on homeowners (§2.6). (In the jargon: the orderly markdown to fundamentals is a *pecuniary externality*, but a benign one; the *constrained-inefficient* loss a planner would want to correct is the synchronized fire-sale overshoot that binding covenants amplify, which the building-by-building stagger does correct, Dávila and Korinek 2018.)

Proposition 4 (Prudent-use assessment stabilizes the base)**.** Valuing space idle past a threshold at stabilized prudent-use NOI—capitalized at a *normalized* rather than a blown-out current cap rate, which would otherwise re-import the spillover—instead of in-place distressed NOI makes the covenant-breach penalty independent of neighbors' distress, which (i) removes the contamination stake that deters a unilateral cut, so the freeze is no longer self-enforcing (though it does not by itself force the most-levered to act—that is the escalating tax's job); (ii) raises the idle owner's own ad valorem carry slightly, by assessing above distressed value, so it mildly reinforces the stick rather than working against it; and (iii) keeps the commercial base on the rolls as transactions print low comps, so the transition does not redistribute the levy onto homeowners (§2.6). It is the base-stabilizing complement to the escalating tax, not a substitute for the forcing function.

4 Related Work

The closest antecedent is Stackman and Moszkowski. *Bleaker on Broadway* (Stackman and Moszkowski 2022) shows below-market-lease covenants impose rent floors and materially raise vacancy at the most-constrained Manhattan storefronts—but it is a *single* bank–landlord pair with *exogenous* market rents: no cross-property comp channel, no strategic interaction. Their companion *Option Value and Storefront Vacancy* (Moszkowski and Stackman 2024) finds a 1%-of-value tax near welfare-neutral and backs out a required externality it leaves unmodeled. We endogenize the market rents *Bleaker* takes as given and supply the externality *Option Value* leaves free. The welfare skeleton is the pecuniary/collateral-externality literature (Lorenzoni 2008; Bianchi 2011; Dávila and Korinek 2018); the nearest empirical cousin is residential foreclosure price-spillovers (Campbell, Giglio, and Pathak 2011). Dynamics build on real-options exercise games (Dixit and Pindyck 1994; Grenadier 1996, 2002) with a sign flip to a war of attrition; default à la Leland (1994); global-games selection (Carlsson and van Damme 1993; Morris and Shin 1998). On the lender side the mechanism mirrors "extend-and-pretend" (Crosignani and Prazad 2024); distress magnitudes are in Jiang et al. (2025) and the structural revaluation in Gupta et al. (2026).

5 Calibration

Public CMBS and assessor data put downtown office leverage in a wide 111–159% range—heterogeneous enough that no single flat tax fits it, which is the case for escalation (§6).

We triangulate current LTV across three free public sources—CMBS servicer filings (SEC Form 10-D), public-REIT 10-Qs, and county assessor records (Table 3).

Table 3: Leverage anchors from public data.

Asset	Source	Current LTV	Role
US Bank Centre	sale + county	153%	high-LTV anchor
Selig 7-building loan	SEC 10-D	∼159%	high-LTV anchor
Selig Office Portfolio (9-bldg)	KBRA / CMBS	∼111%	*counter-anchor*
1201 Third Ave	county	—	re-rate corroboration
Hudson Pacific (portfolio)	SEC 10-Q	—	context (not Seattle-only)

Two high-LTV anchors (∼153%, ∼159%) and one ∼111% counter-anchor (two trace to Selig, so independence is partial), but the *range*—111–159%—is the point: downtown office leverage is heterogeneous. At the high end the most-levered owners are already below the default-trigger rent, so any positive holding tax tips their wait into default rather than a lease (Prop. 3): a flat excise large enough to move them produces a foreclosure, not a tenant. That is precisely why the instrument must *escalate* (§6). The 2024–2026 record matches the prediction: a forced action (Selig's defaults and special servicing) printed low comps (the 9-building pool re-appraised $544.5M→$341.2M—a CMBS appraisal, distinct from the assessor values in Table 1; the US Bank Centre sale confirming the same on the open market), assessors re-rated the block (Table 1), and lower values breached neighbors' covenants (office CMBS delinquency ∼12% by Q1 2026).

6 What Seattle Should Do: A Playbook

Break the hold with a carrot and a stick: an idle-time tax that escalates *on each building's own clock (so waiting becomes terminally unaffordable, and exits stagger instead of synchronizing), prudent-use assessment to hold the base on the rolls through the transition, and cheaper conversion plus new-business help as the productive exit. The target is market-priced tenancy, not the rescue of over-levered owners.*

The package is a carrot-and-stick (Figure 5), aimed at office, where vacancy is ∼36%; downtown retail, at ∼4%, rarely trips the idle-time trigger.

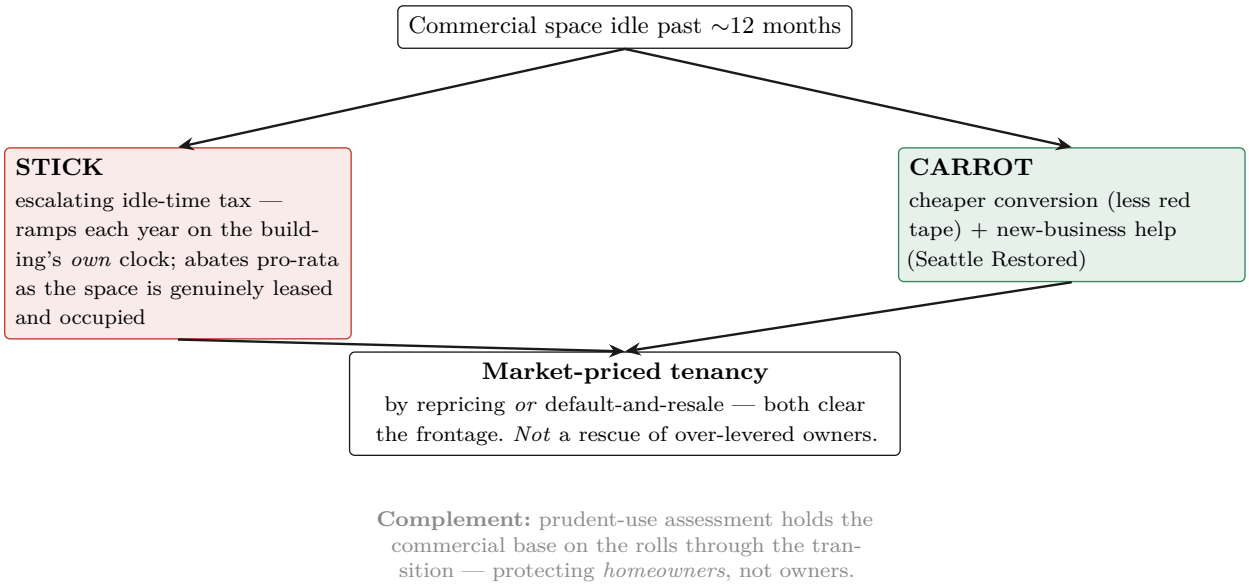

Figure 5: The playbook as a carrot-and-stick: an escalating idle-time tax on each building's own clock, a conversion and new-tenant carrot, and prudent-use assessment as the base-stabilizing complement. Qualified *physical* occupancy—real beneficial use, not a posted or signed rent—switches off the tax; both exits, repricing and default-and-resale, end in market-priced tenancy.

Lead instrument—an escalating idle-time tax. Charge a per-square-foot excise on commercial space vacant past a grace window (∼12 months) and *escalate it each idle year* (roughly $0 in the grace year, then $5→$11→$18→$26/sf/yr across years two to five), pre-set and automatic in the ordinance. Two design choices carry the result. First, the clock runs on *each building's own* vacancy age, so distressed comps print building by building, not in one synchronized wave. Second, the trigger is *qualified physical occupancy*—not a signed lease, not a posted rent. Abatement keys on space in genuine beneficial use (badged employees, utility and access records, an occupancy affidavit) above a square-foot-weighted threshold, charged pro-rata on the dark balance, so partial lease-up is rewarded proportionally. Keying on *use* rather than a signed lease defeats the standard dodge—a friendly master lease to an affiliate

or shell that signs at a committed face and never moves in—and qualifying tenancies must be arm's-length (no entity above a common-ownership threshold; rent recourse to a real operating business). The clock tracks *cumulative* idle months and claws back pro-rata if occupancy lapses, so neither a sham lease nor a revolving low-commitment tenancy can reset it. Charging for dark space rather than a posted rent neutralizes the face-rent-plus-concessions dodge (§2.3). Why escalation works where a flat tax fails (Prop. 3): the freeze is a waiting game, but not a fair one—every owner's clock is set by their own debt, and the most-levered runs out first. A flat tax shortens every clock equally and changes nothing; a tax that escalates on each building's own clock gives the wait a deadline no one can outlast. Each owner is pushed into repricing (the lighter-levered) or default-and-resale (the over-levered), both of which restore market pricing. (Escalation does not by itself make an owner already past the default-trigger rent *lease* rather than default—Prop. 3 still holds; it adds a finite, *staggered* exit date, which is what the frontage needs. Crucially, the tax runs with the parcel through *any* transfer—default, related-party sale, or arm's-length sale to a clean buyer—and never resets on a transaction; it abates only as the space is actually occupied. A buyer therefore prices the accrued, still-escalating charge into what they pay (forcing an honest, low purchase price) and inherits a sharp reason to fill the space fast. This closes the sale-to-reset loophole, turns every transfer into pressure to lease rather than a fresh clock, and does not penalize the clean-basis buyer the workout depends on—they buy the tax burden at a discount, not on top of the price.) Legally the tax is best defended as an *excise* on a voluntary act (keeping leasable space off the market despite a posted rent), not a demand on mere ownership; it compels no tenancy, so it is no rent ceiling or taking (*Quinn v. State* 2023; *High Tide Seafoods*), and as a new municipal excise it would need state enabling authority. Real analogues escalate similarly—San Francisco's storefront vacancy tax ($250/$500/$1,000 per frontage-foot of retail, years 1/2/3+), D.C.'s Vacant-to-Vibrant Act (an escalating vacant-property rate now phasing in), France's duration-graduated tax (Segú 2020)—while Vancouver, which walked back a pre-announced hike, is the cautionary case that the ramp must be statutory.

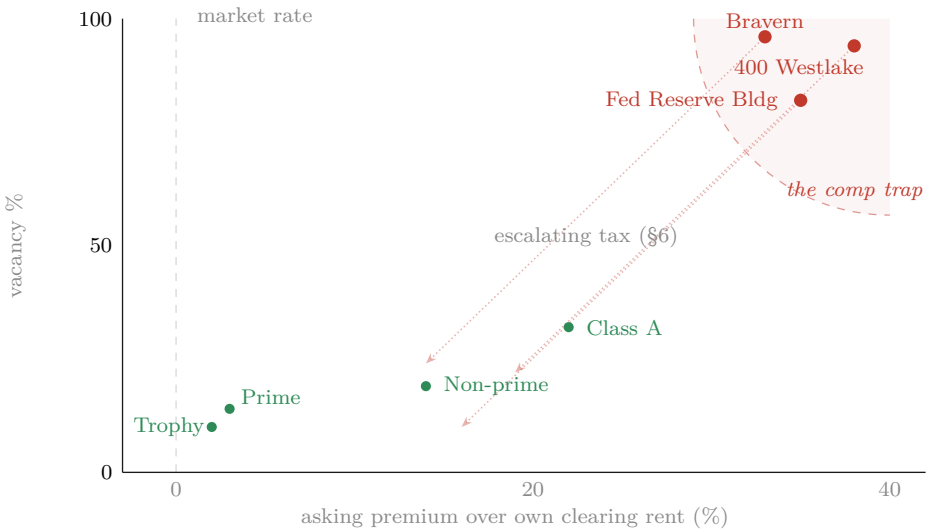

Figure 6: The same axes as Figure 3: each building's vacancy against how far its asking rent sits above its own market-clearing rent. The *comp trap* is the upper-right corner—high asking rent *and* high vacancy, the one combination ordinary quality cannot explain (a desirable building fills; an undesirable one is cheap; only a held comp keeps a building both dear and empty). The escalating idle-time tax (§6) makes waiting terminally unaffordable, so each still-frozen owner must reprice or hand the keys to a clean-basis buyer; either way the building leaves the corner. The arrows trace Seattle's three most-stuck towers—the Bravern, still owner-held, and Selig's 400 Westlake and old Federal Reserve Building, already defaulted onto lenders the tax must follow—moving toward priced-to-clear.

Complement—prudent-use (stabilized-occupancy) assessment. Washington already assesses at highest-and-best-use *market* value, so the reform is not "assess at HBU"; it is to value space idle past the threshold at *stabilized prudent-use NOI*—computed at a fixed long-run submarket occupancy norm and a standard reserve, so an owner cannot re-import distress through the vacancy and reserve assumptions—capitalized at a normalized cap rate (anchored to non-distressed sales of comparable stabilized stock rather than a trailing average that bakes in the fire sales, and fixed in the ordinance rather than left to appraiser discretion), rather than in-place distressed NOI, on the theory that prolonged strategic vacancy is *owner-chosen* economic obsolescence (RCW 84.40.030's "capitalization of income that would be derived from prudent use," with the zoning-use ceiling as the genuine-obsolescence safety valve). This deliberately assesses idle space above its current distressed value, so it is a contested reform under the true-and-fair ceiling, with litigation risk, best drawn as a rebuttable presumption keyed to a defined occupancy threshold. Its job is not to force behavior—the escalating tax does that—but to keep the commercial base on the rolls as low comps print, so the transition does not redistribute the levy onto homeowners (§2.6). If enjoined, the package degrades gracefully: the tax and carrots still force the repricing.

Carrot—cheaper exits and new tenants. Make the productive exit cheap. Conversion is not just an owner's exit; it is the repair for the mono-use zoning that made the shock so severe (§2.1)—every office floor turned to housing puts a resident back in a downtown that had almost none. Seattle has already deferred its sales/use tax on office-to-residential conversions (2025) and waived design review and several development standards; extend those automatically to buildings past the idle threshold and to non-residential reuse, funded from the tax proceeds. On the demand side, scale Seattle Restored—the city's equity-focused activation program (∼190 participants, ∼94% woman- or BIPOC-owned, a dozen pop-ups converting to leases)—so cleared and converted space refills at real, market-clearing rents rather than going dark. Every floor converted out of the office stock also shrinks the comp pool the trap runs on.

The "Bellevue flight" objection. A business-improvement district warns a tax will push tenants to Bellevue, but the relocation literature finds the effect real yet second-order (Giroud and Rauh 2019; Suárez Serrato and Zidar 2016), and agglomeration dampens it further (Brülhart, Jametti, and Schmidheiny 2012); the tax and assessment fall on *owners*, not tenants (Rolheiser 2017); and the documented Seattle→Bellevue drivers are demand-side (Amazon, newer Class A at double the rent).

7 What 2025 Shows: The Break Beginning

The most-levered owner is being forced first, building by building, exactly as the model predicts; and where a market has already taken the hit (San Francisco), price discovery unlocked the recovery.

The freeze is already cracking, exactly where the model says it should. The Selig sequence (§2.4) is the war of attrition resolving in real time: downtown Seattle's most-levered large owner, forced first—not all at once but *building by building* as each loan hit its maturity wall—with the forced sales printing the low comps that then re-rated the whole block (∼53%, about $5.6B, off the top-twenty towers). Staggered over roughly eighteen months and sorted by leverage: that shape is the model's prediction, not a synchronized collapse.

And the break is the cure, once the building changes hands. San Francisco ran the same play about eighteen months ahead. Once a single trophy sale repriced the market at roughly −70% in 2023, price discovery returned—clean-basis buyers could underwrite again, and leasing rebounded to its *highest since 2019* by early 2025, with vacancy still near a third and asking rents still far below peak. The longer the comp is suppressed, the longer the recovery is deferred. So the policy of §6 does not *start* the break—2025 shows it already underway—it *staggers and accelerates* it, so the block reprices building by building rather than in one dump, and the dark frontage returns to use years sooner.

Is it just forbearance? Deepening distress should also drive lenders to renegotiate the terms that generate the floor, and it has: CRE loan modifications rose 66% year-over-year ($16.7B→$27.7B) to mid-2025, 86% in securitized credit, the modal modification a maturity extension, and the 2023 Interagency Policy Statement lets a modified loan keep its *original* appraisal—the value kept off both the rent roll and the balance sheet. But extend-and-pretend only *delays* the comp; it does not generate the cross-sectional pattern that separates the comp trap from simple forbearance (§8).

8 The Test, in One Sentence

Our test keys not on the vacancy level (which demand mostly owns) but on a comparison only the comp trap predicts: when a neighbor falls into distress, a tightly-covenanted owner holds asking rent higher *in response than an otherwise-identical owner who is not.*

The mechanism is already partly measured. The covenant→non-leasing link at the core of the trap is not hypothetical: Stackman and Moszkowski (2022) measure it quasi-experimentally. What is untested is the *cross-building* propagation—that one owner's distress contaminates a neighbor's comp set. The cross-section already shows the predicted sorting (§2.5), and 2025 distress is already driving lenders to renegotiate the very terms that generate the floor (§7).

Take two otherwise-identical buildings—same quality, same block, same quarter—where one owner carries a tightly-covenanted loan and the other does not. When distress hits the *neighbors'* comps, the comp trap predicts the tightly-covenanted owner holds its asking rent *higher* in response. That difference-in-the-response—not the vacancy level (which demand owns), and not even the average held rent—is the falsifiable claim: the paper is wrong if it is zero, if a low comp does not in fact mark down comparable neighbors, or if a concentrated CMBS pool does not enforce the freeze.

Two honest caveats keep the claim sharp. First, the prediction is *signed* only for CMBS pools, where a low comp contractually forces a pool-wide markdown; for a discretionary portfolio bank, which can absorb the loss quietly instead, we leave the sign open, and the test's power comes from the CMBS-versus-bank contrast. Second, the comp trap is *one* testable propagation channel, not the sole cause of held rents—capital starvation, lender-imposed floors, and forbearance can each hold rents up too, and only the interaction above separates the trap from them.

The full specification—the wedge measure, the cross-partial regression, the lender split, the event-study identification, the placebos, and the public-data falsification thresholds—is in Appendix B.

9 Conclusion

A real, structural demand shock made downtown Seattle empty and keeps it empty; we concede that. What demand cannot explain is owners holding *posted* rents high and space dark for years rather than cutting to clear—the self-enforcing equilibrium of a coordination game in which the held *face* rent both guards the owner's collateral and, because tenants finance against it, keeps the space empty. The 2025 record shows that equilibrium beginning to break, most-levered owner first, exactly as the model predicts; and San Francisco shows the break can cure—where a building changed hands at a clean price and the comp finally printed, leasing came back; where the keys went only to a lender holding the line, as in Selig's case, it did not, which is why policy must force the change of hands. The fix is pro-market and modest in ambition: an escalating idle-time tax on each building's own clock, paired with prudent-use assessment and cheaper conversion, that *staggers* a repricing already underway rather than starting one, so the frontage

returns to use years sooner. The mechanism is propagation on a real demand shock; the cleanest test (§8) is the held-rent wedge, not the vacancy level.

A The formal model

This appendix collects the formal environment and the derivations behind the propositions stated in prose in §3.

A.1 Environment

A block has N landlords, each with one vacant space, sharing one comp set. Market rent R_t follows a geometric Brownian motion (drift μ, vol σ, discount ρ, payout yield $\delta = \rho - \mu$), and each building is valued by the income approach, $V_i = \mathrm{NOI}_i/\kappa$. Owner i carries a mortgage $B_i = \theta_0 V_{\text{peak}}$ struck at origination LTV θ_0, with a covenant max LTV θ breached when $V_i < B_i/\theta$; the breach costs a flow $\Pi_i(R) = \lambda \max\{0, B_i/\theta - R/\kappa\}$. The comp channel is $\widehat{R}_i = g(R, \bar{a}_{-i})$ with $\partial g/\partial \bar{a}_{-i} < 0$, so each distressed neighbor lowers V_i and raises Π_i. Marks are *event-driven*: the covenant mark is taken at refinancing, maturity, or transfer to special servicing, not continuously, so in the all-hold state $\Pi_i = 0$ is sustained even though paper values are stale—the markdown does not exist until an owner is forced to a marking event and prints the first low comp. We read R_t crossing the covenant threshold as that event arriving, which is why the maturity wall is the trigger. *Sign-flip:* in Grenadier's exercise game owners race to preempt; here a cheap lease is a public bad, so holding is a strategic complement and "all hold" is a coordination equilibrium—the inversion to a war of attrition governs only *who* is forced to concede first.

A.2 The contamination stake

Write owner i's appraised rent $\widehat{R}_i = g(R, \bar{a}_{-i})$ with *own-comp sensitivity* $g_{\text{own}} \equiv |\partial \widehat{R}_i/\partial a_i|$ and *spillover sensitivity* $g_{\text{nbr}} \equiv |\partial \widehat{R}_i/\partial \bar{a}_{-i}|$; a forced neighbor transaction is discounted by $\chi \in [0, 1)$ as non-representative (a fire sale). **Own-comp dominance** is $g_{\text{own}} > \chi\, g_{\text{nbr}}$. With G_i the lease-up gain from filling now (the rent captured), the contamination stake is

$$S_i = \underbrace{\tfrac{\lambda}{\kappa}\, g_{\text{own}}\, \Delta}_{\text{your own fresh comp}} - \underbrace{\tfrac{\lambda}{\kappa}\, \chi\, g_{\text{nbr}}\, \Delta}_{\text{a neighbor's discounted comp}} + \Omega_i,$$

the covenant damage the *first* low comp inflicts on whoever prints it, net of the smaller damage that owner would suffer if a neighbor moved instead, plus the option value Ω_i spent by cutting (Δ is the size of the cut). Under own-comp dominance, $S_i > 0$.

A.3 Default and the vacancy tax

A levered owner can default rather than lease into a breach. The three exits pay

$$U_{\text{lease}} = \max\{0, R/\kappa - B\} - \Pi, \quad U_{\text{default}} = \max\{0, (1-\phi)R/\kappa - B\}, \quad U_{\text{wait}} = -(m + \tau),$$

where ϕ is the foreclosure haircut, m the holding cost, and τ the vacancy tax. **The tax enters only U_{wait}.** For the deeply underwater anchors $(1-\phi)R/\kappa < B$, so the equity terms clamp at 0; DEFAULT then pays 0 while LEASE still bears the breach flow, $-\Pi < 0$, so default strictly dominates and the deficiency falls on the lender.

A.4 Proof sketches

Lemma (Strategic complements, supporting Prop. 1). Holding preserves the option to lease into an uncontaminated comp set; $\partial^2 W_i/\partial a_i\, \partial(1 - \bar{a}_{-i}) > 0$ when $\lambda > 0$, so the game is supermodular (Topkis). The flow stage is dominance-solvable, but the supermodular *continuation* game admits multiple equilibria—a high-vacancy freeze and a low-vacancy clearing outcome, with global-games selection picking between them. **Prop. 1 (Freeze).** Own-comp dominance $g_{\text{own}} > \chi g_{\text{nbr}}$ gives $S > 0$; the no-unilateral-cut condition $S > G$ makes "all hold" a Nash

equilibrium (the freshest comp on i is i's own, so a unilateral cut breaches i first and by more than the lease-up gain G). Conditional on the freeze, the break runs through the default boundary, not a vanishing stake: with the piecewise-linear penalty, S is constant once breached, so it does not shrink as R_t falls. Since R_{def} rises with leverage (Prop. 2), the most-levered owner is first to reach $R < R_{\text{def}}$, where default dominates holding; it concedes by defaulting. **Prop. 2 (Default).** In the positive-equity region $U_{\text{lease}} = U_{\text{default}}$ gives $\frac{R}{\kappa}(\phi + \lambda) = \frac{\lambda B}{\theta}$, hence $R_{\text{def}} = \kappa\lambda B/[\theta(\phi + \lambda)]$ with $\partial R_{\text{def}}/\partial\theta_0 > 0$. Below zero equity DEFAULT pays 0 while LEASE pays $-\Pi$, so default *strictly* dominates via the avoided Π flow. **Prop. 3 (Policy trap).** Immediate from τ entering only U_{wait}: it cannot change the lease-vs-default ranking, only the wait's attractiveness. **Prop. 4 (Prudent-use).** Stabilized-NOI assessment sets $\partial\widehat{R}_i/\partial\bar{a}_{-i} = 0 \Rightarrow S = 0$; assessing idle space above its distressed value raises that owner's own ad valorem carry (mildly reinforcing the stick), while holding the commercial aggregate on the rolls protects homeowners' share (§2.6).

Limitations. The comp map g is exogenous; its microfoundation is that comp databases observe *face* rents while concessions stay private, so the cross-building signal is face-dominated even though any single appraisal nets concessions—were effective rents fully observable block-wide, S would shrink. We discuss the lender as a coordinating node (§3.2) and the tenant-financing friction (§3.3) but formalize neither, and the friction's vacancy-*level* increment is argued rather than tested—§8's battery isolates the covenant-comp wedge, not the demand-side channel. Mezzanine/preferred standoffs and the tenant-improvement funding constraint likewise stay outside the game—each can hold rents up without the comp channel, and only the §8 cross-partial separates them from it. The freeze is per-owner; default is frictionless; E_{frontage} is asserted, not estimated; the calibration's S and default schedule are illustrative, with the leverage-driven *shape* and the *heterogeneity* the durable core.

B Empirical specification

This appendix gives the formal version of the test stated plainly in §8.

The held-rent wedge. For building i in submarket m at quarter t, with asking rent $A_{i,t}$ and a hedonically-adjusted submarket effective (clearing) rent $\widehat{R}^{\text{eff}}_{m,t}$ (CompStak/CoStar net of free rent and TI), define $\text{Wedge}_{i,t} = \ln A_{i,t} - \ln \widehat{R}^{\text{eff}}_{m,t}$. A positive, *held* wedge is the comp trap's footprint.

The cross-partial (primary test).

$$\text{Wedge}_{i,t} = \beta_1 \,\text{CompDistress}_{i,t} + \beta_2 \,\text{CovTight}_{i,t} + \underbrace{\beta_3 \,[\text{CompDistress} \times \text{CovTight}]_{i,t}}_{\text{comp-trap parameter}} + \gamma' \mathbf{X}_{i,t} + \alpha_i + \delta_{m,t} + \varepsilon_{i,t},$$

with building FE α_i and class×submarket×quarter FE $\delta_{m,t}$ (which absorbs the common demand shock *and* any quality-tier-specific response to it, so β_3 is identified within quality tier). CompDistress is a distance- and class-weighted index of distress (special servicing, delinquency, appraisal-reduction amounts, distressed sales) among i's comparables; CovTight is built from a *pre-distress* baseline—origination LTV/DSCR cushion at the loan's vintage and maturity-wall proximity, not contemporaneous headroom—so the moderator is not itself an outcome of the comp distress it interacts with. **Prediction:** $\beta_3 > 0$**.** The level effect β_1 is contaminated by demand and forbearance; the *cross-partial* is the test. A companion Fine–Gray competing-risks hazard predicts the lease-up hazard falls fastest for tightly-covenanted owners in distressed comp sets ($\theta_3 < 0$), with sale/foreclosure exits signed separately.

The lender discriminant. The sharp split is whether the lender structure *forces* the comp to surface. In a *CMBS* pool a low comp triggers a contractual appraisal reduction that subordinates the junior bonds the special servicer's first-loss affiliate holds; the servicer's duty is to protect pool recovery, so concentration *enforces* the freeze ($\beta_4^{\text{CMBS}} > 0$, rising with pool concentration—a pattern blanket forbearance does not produce). A discretionary portfolio *bank* is ambiguous: it can absorb the loss quietly (breaking the freeze) or extend-and-pretend (§7), so we do not sign β_4^{bank}; the power comes from the CMBS leg and the CMBS-vs-bank contrast. We interact the cross-partial with the lender's comp-set Herfindahl and lender type.

Competing explanations. Held asking rents amid vacancy are over-determined: an owner may post above clearing because it cannot fund tenant-improvement and leasing outlays; because a lender's major-lease-approval right sets the floor; because it is surrendering the building; because a mezzanine or preferred holder blocks a loss-crystallizing lease; or because the lender is forbearing (§7). The all-equity, trophy, and REIT placebos separate the comp trap from demand and mark-to-market but not cleanly from capital-starvation and forbearance, which also switch off without leverage. Quality is the subtler threat—"stranded" low-grade stock (§2.1) was financed hardest at the peak and is emptiest on demand alone—which the class×submarket×quarter effects absorb. What isolates the comp trap is the *interaction*: the wedge's sensitivity to a *neighbor's* distress, rising in covenant tightness and (for CMBS pools) in lender concentration.

Identification and falsification. Leverage is endogenous; we use (i) plausibly-exogenous variation from *origination vintage*, the maturity wall, and CMBS-vs-portfolio assignment; (ii) a stacked difference-in-differences *event study* (Callaway–Sant'Anna) around distressed sales and CMBS appraisal-reduction events, testing whether a low comp marks down *comparable neighbors* and stiffens their asking rents versus matched controls (parallel pre-trends are the falsifiable backbone); and (iii) the lender-concentration heterogeneity. A public-only MVP (SEC EDGAR 10-D/8-K, King County recorder and Assessor extracts, CoStar) supports all three; CompStak upgrades the wedge to a directly-observed measure. **The paper is falsified if** $\beta_3 \approx 0$, if the event study shows flat post-event coefficients or non-parallel pre-trends, or if CMBS-pool concentration does not enforce the freeze ($\beta_4^{\text{CMBS}} \leq 0$). We pre-register $\beta_3 > 0$, $\theta_3 < 0$, $\beta_4^{\text{CMBS}} > 0$, parallel pre-trends, and the placebo nulls (the effect vanishing for all-equity, trophy, and REIT owners). Seattle alone is underpowered for the triple interaction, so we pool several distressed CBDs (San Francisco, Denver, Chicago) with city×quarter effects.

References (selected)

Bianchi (2011), "Overborrowing and Systemic Externalities in the Business Cycle", *AER* 101(7):3400–3426.

Brülhart, Jametti & Schmidheiny (2012), "Do Agglomeration Economies Reduce the Sensitivity of Firm Location to Tax Differentials?", *Economic Journal* 122(563):1069–1093.

Campbell, Giglio & Pathak (2011), "Forced Sales and House Prices", *AER* 101(5):2108–2131.

Carlsson & van Damme (1993), "Global Games and Equilibrium Selection", *Econometrica* 61(5):989–1018.

Crosignani & Prazad (2024), "Extend-and-Pretend in the U.S. CRE Market", FRBNY Staff Reports 1130.

Dávila & Korinek (2018), "Pecuniary Externalities in Economies with Financial Frictions", *Rev. Econ. Stud.* 85(1):352–395.

Dixit & Pindyck (1994), *Investment Under Uncertainty*, Princeton Univ. Press.

Forouhar, Chapple, Allen, Jeong & Greenberg (2025), "Assessing Downtown Recovery Rates and Determinants in North American Cities after the COVID-19 Pandemic", *Urban Studies* 62(6):1209–1231.

Giroud & Rauh (2019), "State Taxation and the Reallocation of Business Activity", *JPE* 127(3):1262–1316.

Glaeser, Kallal, Scheinkman & Shleifer (1992), "Growth in Cities", *JPE* 100(6):1126–1152.

Grenadier (1996), "The Strategic Exercise of Options", *J. Finance* 51(5):1653–1679; (2002), "Option Exercise Games", *RFS* 15(3):691–721.

Gupta, Mittal & Van Nieuwerburgh (2026), "Work From Home and the Office Real Estate Apocalypse", *AER* 116(2):674–709 (NBER w30526).

Jiang, Matvos, Piskorski & Seru (2025), "Monetary Tightening, Commercial Real Estate Distress, and US Bank Fragility", *JPE Macro* 3(4):525–573.

Leland (1994), "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure", *J. Finance* 49(4):1213–1252.

Lorenzoni (2008), "Inefficient Credit Booms", *Rev. Econ. Stud.* 75(3):809–833.

Morris & Shin (1998), "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks", *AER* 88(3):587–597.

Stackman & Moszkowski (2022), "Bleaker on Broadway: The Contractual Origins of High-Rent Urban Blight" (working paper); Moszkowski & Stackman (2024), "Option Value and Storefront Vacancy in New York City" (working paper).

Oates (1969), "The Effects of Property Taxes and Local Public Spending on Property Values", *JPE* 77(6):957–971.

Quan & Quigley (1991), "Price Formation and the Appraisal Function in Real Estate Markets", *J. Real Estate Finance & Econ.* 4(2):127–146.

Rolheiser (2017), "Commercial Property Tax Incidence: Evidence from Urban and Suburban Office Rental Markets", SSRN 2993371.

Segú (2020), "The Impact of Taxing Vacancy on Housing Markets: Evidence from France", *J. Public Econ.* 185:104079.

Shleifer & Vishny (1992), "Liquidation Values and Debt Capacity", *J. Finance* 47(4):1343–1366.

Suárez Serrato & Zidar (2016), "Who Benefits from State Corporate Tax Cuts?", *AER* 106(9):2582–2624.

Legal: Quinn v. State, 526 P.3d 1 (Wash. 2023); *High Tide Seafoods v. State*, 106 Wn.2d 695 (1986); RCW 84.40.030, 84.55, 82.32.330.

Vacancy-tax precedents: San Francisco Commercial Vacancy Tax (Prop. D, 2020); Washington, D.C. Vacant to Vibrant Amendment Act (D.C. Law 26-41, 2025); Vancouver Empty Homes Tax (2017–); Oakland Measure W (2018).