

Conditioned Emptiness: Realized Use as the Correct Denominator for Embodied Burden

Flyxion

Independent Researcher

September 2026

Abstract

An asset can be efficient by every operational measure and still be waste if almost no one uses it, but binary vacancy statistics and per-area intensity metrics both hold capacity fixed in the denominator and so cannot register underuse within occupied or leased stock. This paper defines realized use as capacity-equivalent service hours, $U(K, \tau)$, rather than a binary occupied/vacant flag, and a use-normalized intensity, $I_{\text{realized}} = M_{\text{total}}/U$, temporally and dimensionally consistent with the burden accounting of a companion paper, *Manufactured Necessity*. It keeps three conclusions separate that a single ratio tends to collapse: underutilization, avoidable burden, and withheld capability, and it treats a large I_{realized} as a diagnosis rather than a verdict, following it with an explicit comparison among prospective interventions, restored use, consolidation, or conversion, since the lowest-burden response is not always to fill the asset back up. Two worked examples are used: residential units held out of active use in New York and London, and office towers whose occupied stock runs at markedly less than full design attendance.

1 Introduction

An efficient, well-run, nearly empty building is waste in a sense a vacancy rate and a per-square-foot intensity metric both miss, since both hold capacity fixed in the denominator. This paper’s task is to put realized use in the denominator instead, correctly, and to keep the resulting diagnosis separate from the separate question of what to do about it. Section 3 defines realized use as a capacity-equivalent quantity rather than a binary one. Section 4 puts the numerator on a single, temporally consistent basis. Section 5 states what a large I_{realized} does and does not establish. Section 6 compares prospective interventions rather than assuming restored occupancy is the answer. Section 7 extends the system boundary past the asset itself. Sections 8 and 9 apply the apparatus to residential and office worked examples, and Section 10 checks whether the conclusions survive reasonable changes in assumption.

2 Inherited Apparatus

This paper adopts, without rederiving, the apparatus of *Manufactured Necessity*: capital is discretionary relative to a stated service q and substitute A ; burden is expressed on a declared basis S_0 as a scalar M with a residual R for components that do not reduce to that basis; and a case is manufactured necessity when a large avoided burden coincides with negligible capability loss, $\Delta C \leq \varepsilon$. What follows concerns capital that already exists, so the question is not whether to build but how a given regime of use compares with prospective alternatives to it.

3 Measurement of Realized Use

Definition 3.1 (Capacity-equivalent occupied time). *For asset K over $[0, \tau]$, let $n_K(t)$ be realized occupancy at time t and n_K^* the declared service capacity. Define*

$$U(K, \tau) = \int_0^\tau \frac{n_K(t)}{n_K^*} dt, \quad \rho_{use}(K, \tau) = \frac{U(K, \tau)}{\tau}$$

U is capacity-equivalent occupied time; $\rho_{use} \in [0, 1]$ is the realized utilization fraction.

This replaces a binary occupied-hours count, which cannot distinguish one employee entering a thousand-person office for eight hours from all thousand attending: both register as eight occupied hours. U registers the two cases as 8/1000 and 8 capacity-equivalent hours respectively, which is the distinction the paper needs to make.

$n_K(t)$ and n_K^* must be chosen for the service under examination and not interchanged. For an office, they are plausibly workers present and designed workstation capacity. For housing, person-hours, resident-bedroom-hours, and occupied-unit-hours answer different questions, occupancy intensity, crowding, and simple presence respectively, and a reported figure should state which is meant.

Remark 3.2 (Proxy hierarchy). *Direct counts or sampled occupancy sensors are preferable to access-badge records; badge records are preferable to lease or ownership status, which measures nominal possession rather than use; utility consumption is supporting evidence, not a direct population measure, since a system left running unattended can consume without anyone present. Every reported figure should state its proxy, the proxy's likely direction of bias, and whether it measures entry, duration, intensity, or mere possession.*

Remark 3.3 (Three quantities, kept separate). *H_{binary} records only whether an asset was used at all in a period; $U_{capacity}$, as defined above, measures the utilized fraction of declared service capacity; P_{hours} , total person-hours of exposure to the service, measures human reach rather than capacity utilization. This paper's main metric uses $U_{capacity}$. Distributional comparisons, one resident against several, require P_{hours} instead, and the two are not substitutes for each other.*

4 Temporal and Dimensional Consistency

Definition 4.1 (Total burden over an observation period). *On a declared basis S_0 ,*

$$M_{total}(K, \tau) = M_{operation}(K, \tau) + M_{build}(K) \cdot \frac{\tau}{T}$$

where $M_{operation}(K, \tau)$ is operating burden accumulated over τ , and the embodied term is the asset's construction burden prorated by the fraction of its service life T that τ represents, rather than a rate divided by τ .

Definition 4.2 (Realized-use intensity).

$$I_{realized}(K, \tau) = \frac{M_{total}(K, \tau)}{U(K, \tau)}$$

$M_{operation}$ and M_{build} are additive only once both are expressed on the same declared basis S_0 , as Section 3 of the companion paper requires; M_{total} should not silently combine kilowatt-hours, kilograms of material, and monetary cost. Where a single basis would obscure more than it clarifies, report the vector

$$\mathbf{I}_{realized} = (I_E, I_C, I_W, I_L, \dots)$$

one intensity per retained physical unit (energy, carbon, water, land), rather than forcing a weighted scalar; the scalarization rule of the companion paper still applies wherever a single figure is wanted.

Remark 4.3 (Allocative, not causal). *Amortizing embodied burden across realized use shows how much construction burden each unit of present service is being asked to justify; it does not mean that increasing present occupancy recovers past emissions, which are already spent regardless of what happens to U going forward. I_{realized} is a retrospective, allocative statistic. Any claim about what should be done next is a separate, prospective comparison, taken up in Section 6.*

5 Conditioned Emptiness as a Distinct Case

Manufactured necessity, in the companion paper’s sense, compares an asset’s world against a world in which it was never built, and the avoidable quantity is the entire embodied burden along with dependent infrastructure. Here M_{build} is sunk: no decision available now recovers it. The question is instead whether the asset’s current use, and the operating burden that use requires, is justified relative to prospective alternatives available for the same already-built asset.

A large, sustained I_{realized} establishes only that existing burden is poorly utilized. It does not by itself establish that restoring full occupancy is the correct response, since filling an asset back up generally increases operating burden rather than reducing it, and it does not establish which of several possible interventions dominates. What warrants which intervention is a separate counterfactual comparison, taken up next.

6 From Diagnosis to Intervention

Let S_0 be the asset’s current state and consider three alternatives: A_1 , increase active use; A_2 , consolidate occupants into less space and decondition or mothball the remainder; A_3 , convert the asset to a different service. For each,

$$\Delta M_i = M(S_0) - M(S_{A_i}), \quad \Delta C_i = C(S_0) - C(S_{A_i})$$

Definition 6.1 (Preferred intervention).

$$A^* = \arg \max_{A_i} \Delta M_i \quad \text{subject to} \quad \Delta C_i \leq \varepsilon$$

The three alternatives are not interchangeable in general. Reoccupying an asset (A_1) can lower I_{realized} while raising total burden, since more use also means more operating burden. Consolidation (A_2) can reduce operating burden while removing redundancy or accessibility a distributed pattern of use provided. Conversion (A_3) can restore capability at the cost of substantial new material burden of its own. Which alternative maximizes avoided burden subject to an admissible capability loss is an empirical question specific to the asset class, addressed for each worked example below.

7 System Boundary and Displaced Burden

Definition 7.1 (System burden).

$$M_{\text{system}} = M_{\text{asset}} + M_{\text{access}} + M_{\text{substitute}} + M_{\text{duplication}}$$

where M_{access} is travel burden incurred to reach or leave the asset, $M_{\text{substitute}}$ is the burden of whatever service the asset’s absence or reduced use displaces elsewhere, and $M_{\text{duplication}}$ is burden duplicated across the asset and its substitute rather than shared.

Burden can migrate rather than disappear. An office running below design attendance may show a favorable I_{realized} for the building alone while its absent workers heat, cool, and light hundreds of separate homes, duplicating equipment the office already provided; a program that restores attendance may improve the building’s own I_{realized} while adding a commuting burden the remote pattern had removed. A residence held out of use raises a parallel question: whether its nominally absent owner occupies a substitute dwelling whose burden this accounting has not yet counted. The correct unit of analysis is the provision of the service, not the visible building alone.

8 Worked Example I: Residential Units Held Out of Use

8.1 Scale

New York City’s Housing and Vacancy Survey has counted apartments vacant for “seasonal, recreational, or occasional use,” a category dominated by pieds-à-terre and investment holdings, at 74,945 units, 2.1 percent of the city’s housing stock; a separate American Community Survey estimate places second homes in Manhattan alone near 44,000. London’s Assembly Research Unit reports 105,138 empty homes in the city in 2025, up 81 percent since 2016, with a further 55,978 units recorded as second homes; long-term vacancy (over six months) rose 138 percent over the same period. An on-the-ground account of one prime central London square describes 60 to 65 percent of properties held by international owners present no more than two weeks a year, with most of the remainder used one or two nights a week, consistent with Kensington and Chelsea and Camden recording the highest vacancy shares of any London borough.

8.2 Applying the corrected metric

Take n_K^* as one resident-equivalent occupying the unit full time and $n_K(t)$ as actual presence. A unit visited two weeks a year has $\rho_{\text{use}} \approx 2/52 \approx 0.04$; one used one or two nights a week has ρ_{use} in the range of 0.14 to 0.29. Since $M_{\text{build}} \cdot \tau/T$ does not fall with reduced presence, and standby heating, cooling, and security draw continues on an unoccupied unit, I_{realized} for either pattern rises by roughly the inverse of ρ_{use} relative to full-time occupancy: on the order of twenty-five-fold for the two-week case, three- to sevenfold for the weekly case. These figures describe an individual unit under the stated assumptions; no claim is made that they characterize the housing stock as a whole without the sensitivity treatment in Section 10.

8.3 Capacity, not a unit-to-person count

The comparison between roughly 75,000 seasonal-use units and the city’s shelter population is dimensionally loose as a bare count: units are not people. Converting it requires explicit assumptions. At an average household size for the relevant unit type, plausibly one to two persons for a studio or one-bedroom pied-à-terre, and setting aside units unsuitable for conversion (physically substandard, encumbered by ownership structures resistant to reallocation, or too small for a given household), a defensible range is that the seasonal-use stock could plausibly house on the order of seventy-five thousand to one hundred fifty thousand additional residents if fully returned to primary occupancy, a range that comfortably contains, but is not identical to, the city’s reported shelter population at the time the vacancy figure was recorded. This is offered, as in the companion paper’s treatment of opportunity cost, as an illustrated feasible-alternative scenario under stated assumptions, not as a demonstrated one-to-one correspondence between units and people, and not as a claim that legal, financial, or logistical conversion barriers are minor.

8.4 Intervention comparison

For residential units, A_1 (restore active occupancy) is the plausible dominant alternative: unlike an office, a housing unit’s stated service, shelter, is not substitutable by a remote equivalent, so restoring use directly restores the withheld capability without the system-boundary complication Section 7 raises for offices. A_2 (consolidation) has no clear analogue for a single already-subdivided unit, and A_3 (conversion, for example to short-term rental) restores some capability at markedly lower occupancy intensity than full-time residency and is not treated here as equivalent to A_1 .

9 Worked Example II: Office Towers

9.1 Scale

CoStar’s mid-2026 forecast places national U.S. office vacancy near 14 percent, describing only the binary occupied/leased split. Kastle Systems’ badge-access barometer recorded a record post-pandemic weekly average of 56.3 percent attendance relative to a 2020 pre-pandemic baseline in 2026, with a single-day Tuesday peak of 66.0 percent and Monday and Friday running well below the weekly average. Vacancy and attendance diverge by building grade: CBRE’s Q1 2026 data shows overall Midtown Manhattan vacancy near 18.6 percent against 12.7 percent prime-grade vacancy, while Kastle separately reports Class A+ attendance near 79 percent against a broader-stock average near 56 percent.

9.2 Applying the corrected metric

Taking n_K^* as pre-pandemic designed attendance and $n_K(t)$ as badge-recorded presence, ρ_{use} for the broader stock is on the order of 0.56 at its 2026 peak, falling further on Mondays and Fridays, and $M_{\text{build}} \cdot \tau/T$ is unchanged by how many desks are occupied on a given day. The World Green Building Council estimates embodied, upfront carbon will account for roughly half the whole-life carbon footprint of new construction through 2050 as operational energy decarbonizes, so the fixed component of M_{total} is growing as a share of the total precisely as ρ_{use} in much of the existing stock sits near one half.

9.3 Intervention comparison and system boundary

Here A_1 is not clearly dominant. Office attendance’s stated capability, coordinated in-person work, has a demonstrated remote substitute already in place for the population currently working hybrid, so restoring full attendance does not straightforwardly restore a capability that has actually been lost; it restores a pattern of use. Applying Section 7’s system boundary, a return-to-office program that raises ρ_{use} and lowers the building’s own I_{realized} simultaneously adds M_{access} , commuting burden, that a hybrid pattern had reduced, while $M_{\text{substitute}}$, home-office energy and equipment duplication, falls only partially, since many hybrid workers maintain a home setup regardless. A_2 , consolidating occupied floors and deconditioning or subletting the remainder, plausibly dominates A_1 on ΔM for a building whose ρ_{use} sits structurally near one half rather than dipping temporarily, since it reduces $M_{\text{build}} \cdot \tau/T$ ’s effective denominator problem by shrinking the conditioned footprint to match realized demand rather than by adding commuting burden to fill the existing one. Whether A_2 clears the admissible capability threshold, without degrading collaboration or access for the consolidated population, is not established here and would require case-specific evidence this paper does not have.

10 Sensitivity and Uncertainty

Headline multipliers in Sections 8 and 9 are reported as ranges, not points, because I_{realized} is sensitive to service-life assumption, capacity baseline, and occupancy-proxy choice. For the residential case, varying assumed service life T between forty and eighty years and household size between one and two persons per unit moves the seasonal-use multiplier across roughly a fifteen- to thirty-fivefold range rather than the single twenty-five-fold figure used illustratively above; the qualitative conclusion, that a two-week-a-year unit carries an order of magnitude higher realized-use burden than a fully occupied one, holds throughout that range. For the office case, varying the treatment of weekends and holidays, the assumed coverage gap in badge-access data (buildings with alternate entry systems are known to be undercounted), and designed-capacity assumptions moves ρ_{use} for the broader stock across roughly 0.45 to 0.60; the qualitative conclusion that the broader stock runs at materially less than pre-pandemic design attendance, while Class A+ stock runs closer to it, holds throughout that range. Neither conclusion is reported as robust to changes larger than these bounds, since no attempt has been made here to test them.

11 Discussion

The corrected metric keeps three claims separate that the earlier draft’s single ratio tended to collapse: an asset can be underutilized without every unit of that underutilization being avoidable burden, and avoidable burden can exist without any capability currently being withheld from someone who would use it, as the office case illustrates directly. I_{realized} diagnoses the first; Section 6’s intervention comparison addresses the second; $OC(K)$, inherited from the companion paper, addresses the third. None of the three stands in for either of the others, and the worked examples above land in different places along that separation deliberately: the residential case couples underutilization to withheld capability because housing has no distributed remote substitute, while the office case largely does not, because much of its stated capability now has one.

12 Conclusion

Realized use, measured as capacity-equivalent occupied time rather than a binary flag, and burden held to a single declared basis over a consistent observation period, together give a defensible diagnostic: I_{realized} shows where embodied and operating burden is being converted into comparatively little service. That diagnosis does not by itself select an intervention. Restoring occupancy, consolidating it, or converting the asset to another use carry different prospective burden and capability changes, and the system boundary must extend past the asset to the access and substitute burdens its use pattern displaces elsewhere before any of the three can be called an improvement.

References

- London Assembly Research Unit. *Empty Homes in London*, 2025/2026 data release, Mayor of London.
- Ministry of Housing, Communities and Local Government. *Council Taxbase 2025 in England*; Table 615, vacant dwellings by local authority district.
- Regional Plan Association / NYC Housing and Vacancy Survey data, as reported in Brick Underground and City & State New York coverage of pied-à-terre and seasonal-use vacancy.

- Spears (spearswms.com), "Pretty Vacant," reporting on prime central London buy-to-leave patterns.
- CoStar Group. *U.S. Office Vacancy Forecast*, February and April 2026 revisions.
- CBRE Research. *Q1 2026 U.S. Office Figures*, Midtown Manhattan vacancy data.
- Kastle Systems. *Back to Work Barometer*, 2026 badge-access occupancy data, ten-city and Class A+ indices.
- World Green Building Council. *Bringing Embodied Carbon Upfront*, 2019, and subsequent whole-life carbon estimates.