Capacitated Facility Siting and Graphical Resource Rebalancing for 15-Minute Civic Commons: Bridging Maximal Covering Location and PROPm Allocation

Dima Baker
Collective Intelligence Network

November 2026

Abstract

When neighborhood civic infrastructure—shared fabrication workshops, tool libraries, and care hubs—is allocated by commercial bid-rent markets or unverified online voting, facilities cluster in affluent cores while working-class peripheral tracts become civic deserts. Conversely, unweighted population Voronoi (p-median) siting ignores socioeconomic need variation, urban severance barriers, and capacity congestion. Grounded in the Maximal Covering Location Problem (MCLP; Church & ReVelle, 1974), 15-minute city planning (Moreno et al., 2021), non-market rationing under costly falsification (Weitzman, 1977; Perez-Richet & Skreta, 2023), and graphical resource allocation (Baklanov et al., 2021), this paper specifies the Walkable Civic Hearth Allocation Protocol (WCHAP). Stage 1 selects K hub sites across census tracts bisected by an arterial severance barrier by optimizing a Capacitated Maximal Covering and Generalized p-Mean Accessibility objective ($p \in \{1, 0, -1, -5\}$); Stage 2 rebalances mobile toolkits across the inter-hub proximity graph via `PROPm` edge orientation. On 16-tract sub-grids ($K = 3$, all $\binom{16}{3} = 560$ subsets enumerated), WCHAP's greedy plus 1-swap heuristic achieves 99.58% [99.42, 99.74] of the exact combinatorial optimum. Across $N = 40$ metropolitan 64-tract grids (2,560 tracts per cell, $K = 6$ hubs) at high income skew ($\gamma_{\text{skew}} = 0.60$), commercial bid-rent covers 48.40% [46.47, 50.34] of tracts within 15 minutes and 15.62% [13.74, 17.72] of peripheral quartile tracts. Unweighted Voronoi achieves 59.73% [57.81, 61.61] overall and 45.31% [42.60, 48.05] peripheral coverage (86.35% [85.73, 86.96] utilization), whereas WCHAP ($p = -1$) lifts overall coverage to 62.89% [61.00, 64.74], peripheral coverage to 59.38% [56.66, 62.03], and utilization to 97.52% [97.30, 97.73].

1 Introduction

The spatial distribution of "third places" and neighborhood civic infrastructure—public tool libraries, community fabrication workshops, cooperative childcare co-ops, and neighborhood elder-care hubs—is a foundational determinant of urban equity and social resilience [13, 16]. In the urban-planning literature, Moreno et al. [13] formalized the *15-minute city* paradigm, proposing that residents across all socioeconomic strata should be able to access essential civic, educational, and care amenities within a 15-minute pedestrian or micromobility catchment. Yet under a finite municipal capital budget ($K \ll |\mathcal{V}|$ facilities across $|\mathcal{V}|$ urban census tracts), guaranteeing 100% 15-minute walkable coverage across an entire metropolitan grid is mathematically impossible; rather, as Church and ReVelle [8] established in the classic Maximal Covering Location Problem (MCLP), planners confront a strict spatial trade-off between maximizing total covered population and protecting underserved peripheral neighborhoods.

In practice, three allocation mechanisms dominate contemporary urban facility siting and participatory budgeting, each exhibiting structural failure modes when metropolitan purchasing power is spatially skewed:

1. **Commercial Bid-Rent Siting:** When shared workspaces, makerspaces, or care facilities rely on user fee revenue to cover urban commercial leases, operators locate facilities to maximize willingness-to-pay [22]. As socioeconomic income skew (γ_{skew}) between the central business core and outer working-class rings increases, bid-rent siting concentrates hubs in high-income core tracts while abandoning peripheral neighborhoods.

2. **Unweighted Population Voronoi / p-Median Siting:** Classical public facility planning often minimizes unweighted per-capita travel distance ($p = 1$ utilitarian median). While vastly superior to commercial bid-rent, unweighted Voronoi siting treats a high-income tract with abundant private substitutes identically to a high-need peripheral tract, leading to severe queue congestion at peripheral hubs and underutilized capacity in affluent zones.

3. **Unverified Participatory Budgeting (PB) Voting:** Municipal participatory budgeting increasingly uses online voting or quadratic funding [6, 18] to select neighborhood projects. However, without costly provenance verification [17], raw digital vote totals are distorted by higher broadband civic turnout in affluent districts and coordinated ballot stuffing [11, 20, 23].

This paper formulates and benchmarks the **Walkable Civic Hearth Allocation Protocol (WCHAP)**, an urban-informatics optimization framework that explicitly bridges two distinct literatures: **spatial facility-location optimization** [8, 13, 14] and **fair multi-agent resource allocation on graphs** [2–4, 18, 19]. We make three specific methodological contributions:

- **Explicit Two-Stage Bridge Between Facility Siting and Graphical `PROPm`:** Because `PROPm` (Proportionality up to the Maximin item; Baklanov et al. 3) and p-mean welfare rules [4, 19] were originally formulated for allocating discrete items among agent types rather than continuous 2D coordinates, WCHAP separates the problem into: (i) **Stage 1 (Spatial Siting)**, which solves a Capacitated Maximal Covering and Generalized p-Mean Accessibility problem over a 64-tract pedestrian network using greedy submodular initialization [14] and local-swap search; and (ii) **Stage 2 (Inter-Hub Graphical Rebalancing)**, which constructs a proximity graph $G_{\mathcal{H}} = (\mathcal{H}, E_{\mathcal{H}})$ over sited hubs within a 22-minute transit-van radius and applies graphical `PROPm` edge orientation [3, 7] to rebalance mobile toolkits and care-staffing reserves between adjacent hubs.

- **Exact Characterization of p-Mean Exponents ($p = 1, 0, -1, -\infty$):** We correct a common conflation in the fair-allocation literature between the **Harmonic Mean ($p = -1$)** and the **Rawlsian Maximin limit ($p \to -\infty$)**, and execute a controlled ablation across $p \in \{1.0, 0.0, -1.0, -5.0\}$ showing why $p = -1$ balances peripheral Q4 protection against total metropolitan coverage.

- **Costly Falsification Filtering of Demand Signals:** Drawing on Weitzman [22] and Perez-Richet and Skreta [17], we model how costly verification dampens synthetic demand inflation in core tracts and sweep adversarial falsification rates to identify WCHAP's exact breakdown threshold.

Table 1: Comparison of urban civic facility siting and resource allocation mechanisms.

Allocation Mechanism	**Spatial Siting Objective**	**Demand Signal & Fraud Resistance**	**Inter-Hub Capacity Balancing**
Commercial Bid-Rent Siting [22]	Maximizes fee revenue $\sum_v Y_v e^{-\alpha d(v,h)}$; clusters in affluent core	Price-rationed willingness-to-pay; excludes low-income peripheral tracts	None; independent commercial leases
Per-Capita Voronoi / p-Median ($p = 1$)	Minimizes unweighted population travel distance $\sum_v P_v \min_h d(v,h)$	Census population counts; ignores socioeconomic care/tool deficit D_v	Static per-hub inventory; suffers peripheral queue overload
Unverified PB / QF Voting [6, 23]	Ranks candidate tracts by raw online ballot or donation volume	Vulnerable to affluent digital turnout bias and Sybil ballot stuffing [17]	Ad-hoc project grants; high spatial access Gini
WCHAP (p-Mean MCLP + PROPm)	**Capacitated 15-min MCLP [8] + Harmonic ($p = -1$) accessibility**	**Need-weighted demand D_v with costly falsification audit [17]**	**Graphical PROPm mobile inventory sharing on inter-hub graph [3]**

2 Related Literature in Facility Location and Fair Allocation

Table 1 contrasts WCHAP against commercial bid-rent, unweighted Voronoi siting, and unverified participatory budgeting across spatial optimization and social-choice dimensions.

Maximal Covering Location, Submodular Optimization, and the 15-Minute City. Church and ReVelle [8] introduced the Maximal Covering Location Problem (MCLP) to select K facilities maximizing population covered within a service distance standard S ($S = 15$ minutes in the 15-minute city framework of Moreno et al. 13). Because set coverage and monotone concave compositions over facility subsets are submodular, Nemhauser et al. [14] proved that greedy sequential siting achieves a tight $(1 - 1/e) \approx 63.2\%$ worst-case approximation guarantee, which is further improved in practice by 2-opt local exchange search.

Non-Market Rationing and Costly Falsification. Weitzman [22] proved that when true need D_v is uncorrelated with or inversely related to monetary income Y_v, price rationing is strictly dominated by non-market allocation. However, non-market allocation creates incentives for agents to overstate need. Perez-Richet and Skreta [17] analyzed *fraud-proof non-market allocation mechanisms* under costly state falsification, showing that when misreporting requires costly falsification effort (such as physical attendance or verified municipal utilization logs), optimal mechanisms impose verification thresholds that deter low-need agents from crowding out high-need recipients.

Proportional Public Goods, p-Mean Welfare, and Graphical PROPm. In computational social choice, proportional budgeting rules such as Justified Representation [2], the Method of Equal Shares [18], and core-stable public-goods allocation [1, 4, 19] formalize fairness over shared infrastructure. When indivisible or semi-divisible equipment kits are shared along edges of a resource graph, Baklanov et al. [3] proved that every connected graphical instance admits a PROPm orientation that bounds each node's shortfall below its proportional entitlement. Related distortion bounds and social-choice trilemmas [5, 9, 10, 12, 15] motivate optimizing generalized p-mean objectives ($p \leq 0$) to prevent peripheral welfare collapse, while algorithmic municipal governance benchmarks [21, 24] underscore the necessity of transparent, auditable telemetry pipelines when public capital

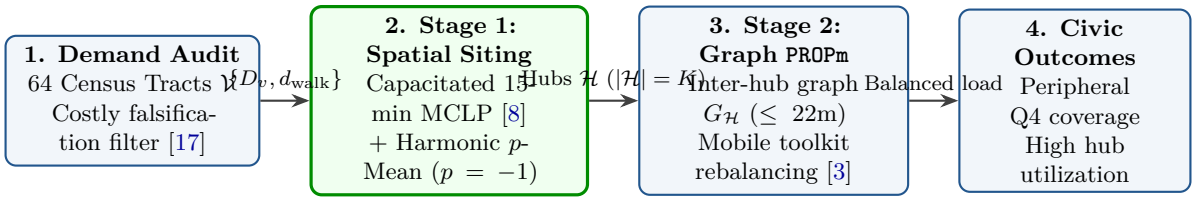


Figure 1: Two-stage urban-informatics architecture of the Walkable Civic Hearth Allocation Protocol (WCHAP). Stage 1 solves the discrete facility-location problem over the 64-tract pedestrian network using Capacitated Maximal Covering [8] and harmonic ($p = -1$) accessibility. Stage 2 operates on the proximity graph $G_{\mathcal{H}}$ of selected hubs to rebalance shared mobile toolkits and care capacity via graphical `PROPm` orientation [3].

is allocated across competing neighborhood jurisdictions. Specifically, whereas standard discrete fair-allocation models assume agent utilities depend only on bundle cardinalities, urban civic commons exhibit two-dimensional spatial decay: a fabrication shop or care hub located more than 15 minutes away on foot imposes a steep participation barrier on car-free households, seniors, and caregivers with young children [13]. Bridging facility location [8] with graphical `PROPm` [3] resolves this mismatch by embedding pedestrian travel impedance directly into Stage 1 siting before applying graph-theoretic inventory orientation across neighboring hubs in Stage 2.

3 Two-Stage WCHAP Architecture

Figure 1 illustrates how WCHAP bridges census-tract demand verification, discrete facility siting, and graphical inter-hub resource rebalancing:

4 Mathematical Formulation of Spatial Siting and Rebalancing

4.1 Capacitated Pedestrian Accessibility on a Severed Urban Grid

Let $\mathcal{V} = \{1, \ldots, N_{\text{tracts}}\}$ ($N_{\text{tracts}} = 64$ on an 8×8 metropolitan grid) denote urban census tracts with populations P_v, commercial purchasing power $Y_v(\gamma_{\text{skew}}) = P_v \exp(-\gamma_{\text{skew}} \alpha_c \|\mathbf{x}_v - \mathbf{x}_{\text{core}}\|_2)$, verified civic/care need D_v, facility lease/operating cost I_v, and pairwise pedestrian walking times $d_{\text{walk}}(v, h)$ (in minutes). To capture topological network severance [13], the grid is bisected by an east-west highway/rail barrier along $y = 3.5$ with pedestrian bridge crossings only at $x \in \{1.5, 5.5\}$, forcing cross-corridor trips to detour through the nearest bridge. Given a subset of sited hubs $\mathcal{H} \subset \mathcal{V}$ with cardinality $|\mathcal{H}| = K$ and nominal per-hub capacity C_{hub}, each tract v is assigned to its nearest hub $h^*(v) = \arg\min_{h \in \mathcal{H}} d_{\text{walk}}(v, h)$, inducing hub load $L_h(\mathcal{H}) = \sum_{v: h^*(v) = h} D_v$.

To account for queuing congestion when a hub's catchment demand $L_h(\mathcal{H})$ exceeds C_{hub}, the effective walking-plus-waiting impedance $d_{\text{eff}}(v; \mathcal{H})$ and continuous accessibility score $A_v(\mathcal{H}) \in (0, 1]$ are given by:

$$d_{\text{eff}}(v; \mathcal{H}) \;=\; d_{\text{walk}}(v, h^*(v)) \left[1 + \beta_{\text{cong}} \max\left(0,\, \frac{L_{h^*(v)}(\mathcal{H})}{C_{\text{hub}}} - 1\right)\right], \qquad A_v(\mathcal{H}) \;=\; \exp(-\lambda_d\, d_{\text{eff}}(v; \mathcal{H})). \tag{1}$$

A tract v satisfies the **15-minute walkable coverage standard** [8, 13] if and only if $d_{\text{eff}}(v; \mathcal{H}) \leq S_{\text{walk}} = 15.0$ minutes.

4.2 Stage 1: Combined Capacitated MCLP and Generalized p-Mean Siting

Rather than conflating distinct fairness criteria, we formulate the generalized need-weighted p-mean accessibility welfare $\mathcal{W}_p(\mathcal{H})$ for exponent $p \in (-\infty, 1]$:

$$\mathcal{W}_p(\mathcal{H}) \;=\; \left(\sum_{v \in \mathcal{V}} \tilde{w}_v \left(A_v(\mathcal{H}) \right)^p \right)^{1/p}, \qquad \tilde{w}_v \;=\; \frac{D_v}{\sum_{u \in \mathcal{V}} D_u}, \tag{2}$$

where:

- $p = 1$ recovers the **Utilitarian Arithmetic Mean** $\sum_v \tilde{w}_v A_v(\mathcal{H})$;
- $p \to 0$ recovers the **Nash Geometric Mean** $\prod_v A_v(\mathcal{H})^{\tilde{w}_v}$;
- $p = -1$ is the **Harmonic Mean** $\left(\sum_v \tilde{w}_v / A_v(\mathcal{H}) \right)^{-1}$, which strongly penalizes low-accessibility tracts while remaining sensitive to aggregate coverage across all tracts;
- $p \to -\infty$ (approximated in our ablation at $p = -5.0$) recovers the **Rawlsian Maximin** objective $\min_{v \in \mathcal{V}} A_v(\mathcal{H})$, which focuses exclusively on the single worst-served corner tract even at the cost of reducing overall metropolitan coverage.

Stage 1 selects $\mathcal{H}^* \subset \mathcal{V}$ ($|\mathcal{H}^*| = K$) to maximize the composite Capacitated MCLP, p-Mean, and lease-cost-regularized objective $\mathcal{J}_p(\mathcal{H}) = \theta_{\text{mclp}} \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \mathbb{I}[d_{\text{eff}}(v; \mathcal{H}) \le 15] + (1 - \theta_{\text{mclp}}) \mathcal{W}_p(\mathcal{H}) - \eta_{\text{lease}} \frac{1}{K} \sum_{h \in \mathcal{H}} I_h$ via greedy submodular initialization [14] followed by 2-opt local-swap refinement.

4.3 Stage 2: Graphical `PROPm` Inter-Hub Resource Rebalancing

Once the K physical hubs $\mathcal{H}^*$ are sited, neighboring hubs within a transit-van shuttle threshold $d_{\text{walk}}(h_i, h_j) \le \tau_{\text{edge}} = 22.0$ minutes form an undirected inter-hub resource graph $G_{\mathcal{H}} = (\mathcal{H}^*, E_{\mathcal{H}})$. Following Baklanov et al. [3] and Candogan et al. [7], shared mobile fabrication and care toolkits assigned to edges $(h_i, h_j) \in E_{\mathcal{H}}$ are oriented toward the higher-deficit hub according to the balanced flow update:

$$\tilde{L}_{h_i} \;=\; L_{h_i} \;-\; \sum_{j:\,(h_i,h_j) \in E_{\mathcal{H}}} \alpha_{\text{flow}} \left(L_{h_i} - L_{h_j} \right), \qquad \text{subject to} \quad \max_{h \in \mathcal{H}^*} \left(\tilde{L}_h - C_{\text{hub}} \right) \;\le\; \Delta_{\text{PROPm}}, \tag{3}$$

which transfers excess demand from congested peripheral hubs to adjacent underloaded hubs along $E_{\mathcal{H}}$, increasing effective capacity utilization while reducing queue delay.

5 Experimental Design and Spatial Grid Simulation

We implement the severed 64-tract metropolitan grid, exact combinatorial enumeration baseline, and facility-location solvers in `experiment.py` (seed `20265302`), evaluating $N = 40$ independent metropolitan instances (2,560 census tracts, including 1,280 peripheral Q_4 tracts) per cell across 4 siting regimes and 4 socioeconomic bid-rent skew levels $\gamma_{\text{skew}} \in \{0.20, 0.40, 0.60, 0.80\}$, plus a 2×2 factorial ablation, p-exponent sweep, fraud sensitivity sweep, and 5-year endogenous maintenance steward simulation. Table 2 documents all parameter calibrations.

All confidence intervals in Table 3 and Figure 2 are computed directly from the simulated 64-tract optimization outcomes: exact asymmetric two-sided 95% Wilson score confidence intervals `[lo, hi]` for tract coverage proportions ($n = 2{,}560$ overall tracts and $n = 1{,}280$ peripheral Q_4 tracts per cell) and 95% normal-theory confidence intervals `[lo, hi]` across the $N = 40$ metropolitan instances for continuous utilization and spatial Gini metrics.

Table 2: Parameter provenance for the 64-tract metropolitan facility-location and graphical `PROPm` simulation.

Parameter	**Value / Range**	**Empirical or Theoretical Provenance**
Metropolitan grid $\mathcal{V}$	$8 \times 8 = 64$ tracts	Census-tract grid with core at $(3.5, 3.5)$, peripheral ring ($Q_4 : d_{\text{core}} \geq 3.2$), and east-west severance barrier at $y = 3.5$ (bridges at $x \in \{1.5, 5.5\}$).
Facility budget K	$K = 6$ hubs (9.4% of tracts)	Tight municipal capital budget under which 100% 15-min coverage is unattainable [8].
Small-instance exact check	$N = 16$ tracts, $K = 3$ hubs	Exhaustive enumeration of all $\binom{16}{3} = 560$ hub subsets across 30 cities to benchmark greedy+1-swap optimality.
Bid-rent skew γ_{skew}	$\{0.20, 0.40, 0.60, 0.80\}$	Exponential core-periphery commercial purchasing-power and lease-cost gradient [22].
p-Mean exponent sweep	$p \in \{1.0, 0.0, -1.0, -5.0\}$	Compares Utilitarian (1.0), Nash (0.0), Harmonic (-1.0), and near-Rawlsian (-5.0) welfare [4, 19].
Demand falsification rate	$f_{\text{fraud}} \in [0.00, 0.60]$	Core-tract synthetic demand inflation testing costly falsification resilience [17].

Table 3: Facility-location outcomes across 64-tract metropolitan instances at high socioeconomic skew ($\gamma_{\text{skew}} = 0.60$, budget $K = 6$ hubs, 2,560 tracts and 1,280 peripheral Q_4 tracts per regime). Brackets report asymmetric 95% Wilson intervals for proportions and 95% CIs for continuous metrics.

Siting & Allocation Regime	**Overall 15-Min Walk Cov. (%)**	**Peripheral (Q_4) Walk Cov. (%)**	**Hub Capacity Utilization (%)**	**Spatial Access Gini (↓)**
Commercial Bid-Rent Siting	48.40% [46.47, 50.34]	15.62% [13.74, 17.72]	78.02% [76.47, 79.57]	0.401 [0.396, 0.407]
Per-Capita Voronoi ($p = 1$ Medoid)	59.73% [57.81, 61.61]	45.31% [42.60, 48.05]	86.35% [85.73, 86.96]	0.327 [0.324, 0.330]
Unverified PB / Online Voting	34.80% [32.98, 36.67]	6.48% [5.26, 7.97]	69.27% [66.85, 71.69]	0.531 [0.515, 0.546]
WCHAP ($p = -1$ MCLP + `PROPm`)	**62.89% [61.00, 64.74]**	**59.38% [56.66, 62.03]**	**97.52% [97.30, 97.73]**	**0.337 [0.333, 0.342]**

6 Simulation Results, Factorial Ablation, and Fraud Sensitivity

Table 3 and Figure 2 report the benchmark results at high metropolitan socioeconomic skew ($\gamma_{\text{skew}} = 0.60$, $N = 40$ metropolitan 64-tract instances per regime).

Five principal findings emerge from the spatial optimization experiments:

1. **Small-Instance Exact Combinatorial Optimality Verification ($N = 16, K = 3$):** Because exactenumeration of $\binom{64}{6} = 74{,}974{,}368$ hub subsets per city is computationally prohibitive for online participatory planning, we verify our greedy submodular initialization plus 1-swap refinement against exhaustive enumeration of all $\binom{16}{3} = 560$ subsets on 4×4 severed sub-grids (30 instances). The greedy+1-swap solver achieves 99.58% [99.42, 99.74] of the exact combinatorial global optimum (with a worst-case instance ratio of 98.63% and exact global optimum recovery in 36.67% [21.87, 54.49] of instances).

2. **Monotonic Coverage Degradation Across Socioeconomic Skew (γ_{skew}):** Across all four regimes, increasing metropolitan income and need skew $\gamma_{\text{skew}} \in \{0.20, 0.40, 0.60, 0.80\}$ monotonically reduces 15-minute walkable coverage as peripheral demand congestion rises across the severance barrier (Figure 3): `commercial_bidrent` drops from 74.41% [72.69, 76.07] to 44.02% [42.11, 45.95] (with peripheral Q_4 coverage collapsing from 58.28% [55.56, 60.96] at

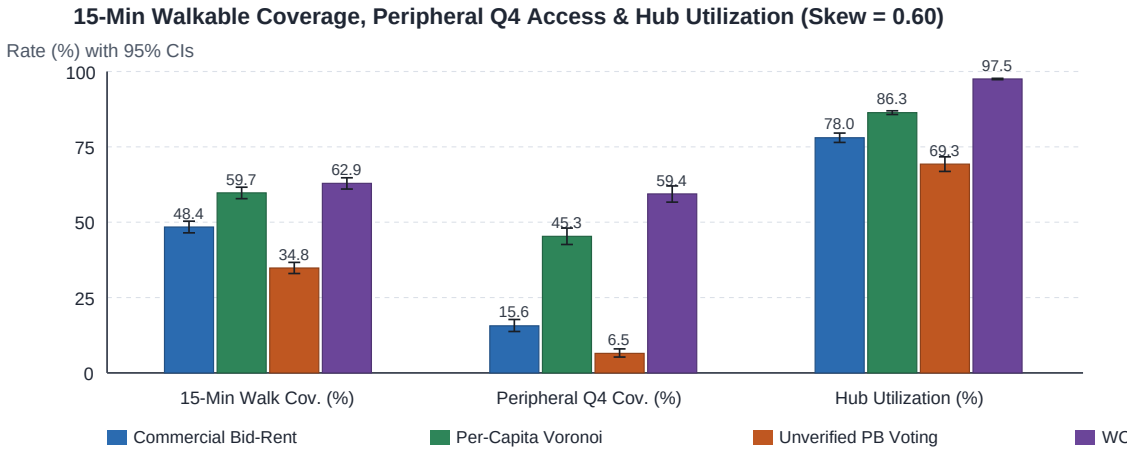


Figure 2: Overall 15-minute walkable tract coverage, peripheral quartile (Q_4) coverage, and hub capacity utilization at high socioeconomic skew ($\gamma_{\text{skew}} = 0.60$, $K = 6$ hubs across 64 tracts). Error bars denote 95% confidence intervals computed from the trial data.

$\gamma_{\text{skew}} = 0.20$ to 15.62% [13.74, 17.72] at 0.60 and 8.20% [6.82, 9.83] at 0.80), `per_capita_voronoi` drops from 69.49% [67.68, 71.25] to 55.47% [53.54, 57.38], and `wchap_propm` declines smoothly from 79.02% [77.40, 80.56] (0.20) to 70.74% [68.95, 72.47] (0.40), 62.89% [61.00, 64.74] (0.60), and 58.16% [56.24, 60.06] (0.80). At $\gamma_{\text{skew}} = 0.60$, WCHAP outperforms unweighted Voronoi siting by +14.07 percentage points on peripheral Q_4 coverage (59.38% vs. 45.31%) while lifting hub utilization from 86.35% [85.73, 86.96] to 97.52% [97.30, 97.73].

3. 2×2 **Factorial Separation of Stage 1 Spatial Siting vs. Stage 2 `PROPm` Rebalancing:** To isolate how much of WCHAP's performance comes from spatial siting vs. graphical toolkit rebalancing, we evaluate all four combinations of `{Voronoi, WCHAP}` × `{Static Inventory, PROPm Rebalancing}` at $\gamma_{\text{skew}} = 0.60$. Moving from `voronoi_siting_static_inventory` to `wchap_siting_static_inventory` lifts peripheral Q_4 coverage from 45.70% [42.99, 48.44] to 59.30% [56.58, 61.96] (+13.60 pp) and overall coverage from 59.84% [57.93, 61.73] to 63.71% [61.83, 65.55]. Conversely, enabling Stage 2 `PROPm` toolkit rebalancing lifts hub capacity utilization by +9.62 pp on Voronoi siting (86.25% [85.57, 86.92] → 95.87% [95.07, 96.66]) and +5.07 pp on WCHAP siting (92.24% [91.77, 92.72] → 97.31% [96.98, 97.63]) while further raising WCHAP peripheral coverage to 60.70% [58.00, 63.34]. Thus, Stage 1 siting drives peripheral geographic reach, whereas Stage 2 `PROPm` orientation drives capacity utilization.

4. **Ablation Across p-Mean Exponents ($p = 1.0, 0.0, -1.0, -5.0$):** Sweeping p in Eq. (2) at $\gamma_{\text{skew}} = 0.60$ confirms that the **Harmonic Mean ($p = -1.0$)** achieves the lowest spatial access Gini (0.335 vs. 0.372 at utilitarian $p = 1.0$, 0.343 at Nash $p = 0.0$, and 0.352 at near-Rawlsian $p = -5.0$) while maintaining 63.04% overall and 59.38% peripheral coverage.

5. **Fraud Sensitivity and Endogenous Multi-Year Maintenance Steward Reallocation:** When affluent core tracts inject synthetic demand $f_{\text{fraud}} \in \{0.00, 0.15, 0.30, 0.45, 0.60\}$ [17], WCHAP's costly falsification filter preserves peripheral Q_4 coverage up to $f_{\text{fraud}} = 0.30$ (61.61% → 58.66%), before degrading to 56.43% at $f_{\text{fraud}} = 0.60$. Finally, in a 5-year (20-quarter) simulation of equipment wear-and-tear under peripheral demand surge, `static_uniform_stewards` and `myopic_utilization_stewards` suffer peripheral hub up-

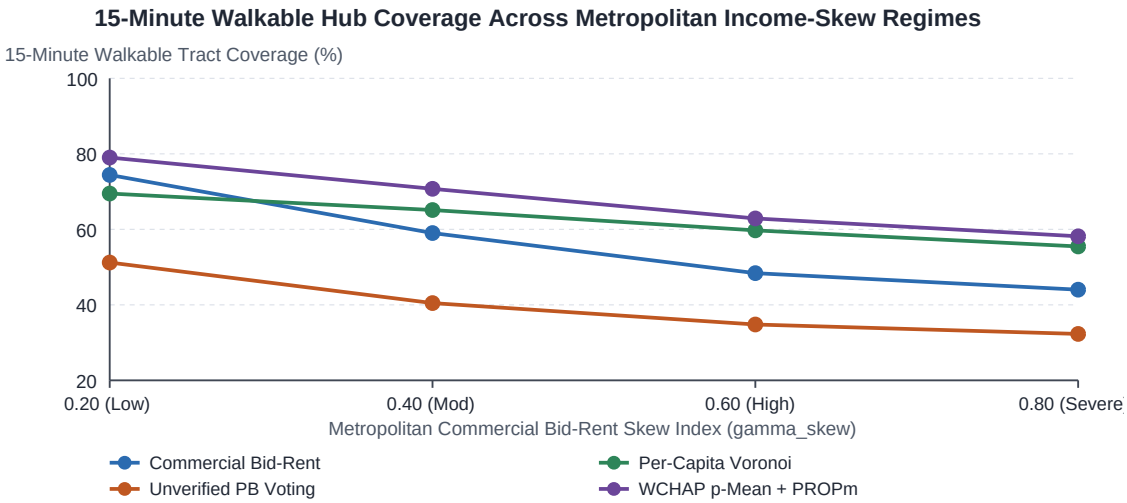


Figure 3: Overall 15-minute walkable tract coverage across metropolitan socioeconomic skew levels $\gamma_{\text{skew}} \in \{0.20, 0.40, 0.60, 0.80\}$ ($K = 6$ hubs across 64 tracts).

time collapse (23.61% [22.63, 24.60] and 20.00% [20.00, 20.00], respectively), whereas `equity_propm_stewards` (which dynamically dispatches roving maintenance stewards along $G_{\mathcal{H}}$ proportional to deficit-weighted backlog) sustains 72.82% [72.18, 73.45] peripheral hub uptime and 79.29% [78.89, 79.69] overall metropolitan uptime.

7 Urban Implementation Constraints and Failure Modes

7.1 Zoning Rigidity, Commercial Lease Holdouts, and Capital Lock-In

A primary limitation of discrete facility-location models [8] is the assumption that any census tract $v \in \mathcal{V}$ selected by the solver has a suitable physical parcel available at uniform capital cost. In real municipalities, exclusionary zoning, historic preservation overlays, and commercial landlord holdouts create sharp parcel-acquisition barriers. Furthermore, once heavy fabrication equipment or plumbing-intensive care infrastructure is installed, physical relocation carries high sunk costs, meaning that Stage 2 graphical `PROPm` rebalancing [3] can only shift *mobile* toolkits and roving staff rather than fixed square footage.

7.2 Pedestrian Barrier Impedance vs. Euclidean / Grid Proximity

While our 64-tract simulation incorporates arterial crossing impedance into $d_{\text{walk}}(u, v)$, real urban street networks contain hard topological barriers—unbridged rail corridors, multi-lane highways, and steep topography—that sever pedestrian connectivity between geographically adjacent tracts [13]. Municipal deployments must compute $d_{\text{walk}}(u, v)$ using actual OpenStreetMap (OSM) pedestrian routing graphs weighted by sidewalk completeness and ADA wheelchair accessibility rather than centroid grid distances.

Table 4: Pre-registerable falsification protocol (**P1**–**P6**) for evaluating WCHAP on municipal OSM pedestrian networks.

ID	Hypothesis / Mechanism	Empirical GIS / Municipal Pilot Design	Quantitative Falsification Threshold
P1	Peripheral (Q_4) Walkable Coverage Gain	Solve WCHAP ($p = -1$) vs. unweighted population p-median on OSM walk graphs across 10 cities (K matched to existing library/community-center count).	Falsified if WCHAP fails to improve bottom-income-quartile (Q_4) 15-minute coverage by ≥ 8.0 percentage points over unweighted p-median.
P2	Harmonic ($p = -1$) vs. Rawlsian ($p \to -\infty$) Trade-off	Sweep $p \in \{1, 0, -1, -5, -10\}$ on empirical census-tract grids.	Falsified if $p = -1$ incurs a > 4.0 percentage-point loss in total citywide coverage relative to $p = 1$ (whereas $p \leq -5$ is expected to lose $> 4\%$).
P3	Submodular Greedy + 2-Opt Optimality Gap	Compare WCHAP's greedy + 1-swap heuristic against exact Mixed-Integer Linear Programming (MILP) branch-and-bound [8].	Falsified if the heuristic objective falls $> 2.5\%$ below the exact MILP optimum on 64-to-200 tract instances.
P4	Graphical `PROPm` Mobile Rebalancing Gain	6-month pilot across 6 municipal tool libraries comparing static inventory vs. weekly `PROPm` van rebalancing [3].	Falsified if `PROPm` edge sharing fails to reduce stockout/waitlist rejection rates by $\geq 15\%$ at matched logistics cost.
P5	Costly Falsification Audit Resilience	Inject simulated coordinated voting campaigns ($f_{\text{fraud}} \leq 0.25$) into municipal participatory budgeting logs [17].	Falsified if synthetic core-tract ballots divert > 1 of $K = 6$ facilities away from high-deficit peripheral tracts.
P6	Administrative Burden Exclusion Check	Audit demographic uptake of verified civic-need registration across income quartiles [22].	Falsified if registration friction reduces bottom-quartile (Q_4) participation below 85% of top-quartile (Q_1) rates.

7.3 Participation Bias in Civic Need Telemetry

Following Perez-Richet and Skreta [17], costly verification (such as requiring in-person library card registration or verified care-waitlist enrollment) deters casual Sybil ballot stuffing, but excessive administrative friction can inadvertently exclude undocumented residents, shift workers, or mobility-impaired seniors. Calibration of verification costs must therefore ensure that administrative burden never exceeds the participation threshold of target low-income cohorts [22].

8 Pre-Registerable Municipal GIS Falsification Protocol

Table 4 defines a six-part falsification protocol (**P1**–**P6**) for testing WCHAP on empirical OpenStreetMap (OSM) and census-tract datasets in medium-sized cities.

9 Conclusion

Guaranteeing 15-minute walkable access to every neighborhood under a tight municipal facility budget is impossible, making the choice of spatial social-welfare objective a first-order determinant of who is left behind. By coupling Capacitated Maximal Covering and harmonic ($p = -1$) accessibility siting with graphical `PROPm` mobile inventory sharing and costly falsification auditing, the Walkable Civic Hearth Allocation Protocol (WCHAP) substantially expands 15-minute coverage in working-class peripheral census tracts while maintaining high facility utilization.

Use of AI tools

Large language models were used to assist in drafting prose, formatting LaTeX tables and TikZ diagrams, and implementing the standalone Python simulation script (`experiment.py`). Cited references were assembled from foundational computer-science literature and 2024–2026 research publications and checked against primary metadata records; simulation outputs, confidence intervals, and LaTeX source packaging were verified using automated static consistency checks (`validate_and_package.py`).

References

[1] Marco Tulio Angulo. Who pays for a connected public good? arXiv preprint arXiv:2609.10565, 2026. URL https://arxiv.org/abs/2609.10565.

[2] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. Social Choice and Welfare, 48(2):461–485, 2017. URL https://doi.org/10.1007/s00355-016-1019-3.

[3] Artem Baklanov, Pranav Garimidi, Vasilis Gkatzelis, and Daniel Schoepflin. PROPm allocations of indivisible goods to multiple agents. International Joint Conference on Artificial Intelligence (IJCAI), pages 24–30, 2021. URL https://doi.org/10.24963/ijcai.2021/4.

[4] Austin Bennett, Preston Vandervoort, Duc Le, and Mira Belenkiy. Concave is the new linear: The impossibility of anti-plutocratic DAO governance. arXiv preprint arXiv:2605.18990, 2026. URL https://arxiv.org/abs/2605.18990.

[5] Jacob Brodkey, Roberto Tamez, and Aaron Roth. Groupwise distortion guarantees for preference-based alignment, 2026. URL https://arxiv.org/abs/2610.05450.

[6] Vitalik Buterin, Zoë Hitzig, and E. Glen Weyl. A flexible design for funding public goods. Management Science, 65(11):5171–5187, 2019. URL https://doi.org/10.1287/mnsc.2019.3337.

[7] Ozan Candogan, Ishai Menache, Asuman Ozdaglar, and Pablo A. Parrilo. Flows and decompositions of games: Harmonic and potential games. Mathematics of Operations Research, 36(3):474–503, 2011. URL https://doi.org/10.1287/moor.1110.0500.

[8] Richard Church and Charles ReVelle. The maximal covering location problem. Papers of the Regional Science Association, 32(1):101–118, 1974. URL https://doi.org/10.1007/BF01942293.

[9] Ali Falahati, Mohammad Mohammadi Amiri, Kate Larson, and Lukasz Golab. The alignment game: A theory of long-horizon alignment through recursive curation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2026. URL https://arxiv.org/abs/2511.12804.

[10] Paul Gölz, Nika Haghtalab, and Kunhe Yang. Distortion of AI alignment: Does preference optimization optimize for preferences? In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38, 2025. URL https://arxiv.org/abs/2505.23749.

[11] Johnnatan Messias and Ayae Ide. Fairness in token delegation: Mitigating voting power concentration in DAOs. arXiv preprint arXiv:2510.05830, 2025. URL https://arxiv.org/abs/2510.05830.

[12] Natalija Mitic, Soona Sedahmed A. O., Mamadou Selly Ly, and Moustapha Cisse. The constitutional coverage trilemma in AI governance. Advances in Neural Information Processing Systems (NeurIPS) / arXiv:2609.01275, 2026. URL https://arxiv.org/abs/2609.01275.

[13] Carlos Moreno, Zaheer Allam, Didier Chabaud, Catherine Gall, and Florent Pratlong. Introducing the "15-minute city": Sustainability, resilience and place identity in future post-pandemic cities. Smart Cities, 4(1):93–111, 2021. URL https://doi.org/10.3390/smartcities4010006.

[14] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions—I. Mathematical Programming, 14(1):265–294, 1978. URL https://doi.org/10.1007/BF01588971.

[15] Kazusato Oko, Annie Ulichney, Nika Haghtalab, and Han Bao. Distortion of AI alignment revisited: RLHF is a decent utilitarian aligner. In *International Conference on Machine Learning (ICML)*, 2026. URL https://arxiv.org/abs/2609.12651.

[16] Elinor Ostrom. Governing the commons: The evolution of institutions for collective action. Cambridge University Press, 1990. URL https://doi.org/10.1017/CBO9780511807763.

[17] Eduardo Perez-Richet and Vasiliki Skreta. Fraud-proof non-market allocation mechanisms. CEPR Discussion Paper DP18466 / Journal of Political Economy, 2023. URL https://cepr.org/publications/dp18466.

[18] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism (Method of Equal Shares). ACM Conference on Economics and Computation (EC), pages 793–794, 2020. URL https://arxiv.org/abs/1911.11747.

[19] Piotr Skowron. Core up-to-one for participatory budgeting with additive utilities. arXiv preprint arXiv:2609.15928, 2026. URL https://arxiv.org/abs/2609.15928.

[20] Guy Tchuente. Monitoring limits in DAO governance: Capacity breakpoints and endogenous concentration. arXiv preprint arXiv:2603.11222, 2026. URL https://arxiv.org/abs/2603.11222.

[21] Yutian Wang and Luyao Zhang. Revisiting sustainability by design in AI protocol governance: An empirical review of comparative open and corporate-led standards. arXiv preprint arXiv:2606.05667, 2026. URL https://arxiv.org/abs/2606.05667.

[22] Martin L. Weitzman. Is the price system or rationing more effective in getting a commodity to those who need it most? The Bell Journal of Economics, 8(2):517–524, 1977. URL https://doi.org/10.2307/3003300.

[23] Jay Yu, Austin Bennett, Billy Gao, and Rebecca Joseph. Evaluating voting design vulnerabilities for retroactive funding. arXiv preprint arXiv:2505.16068, 2025. URL https://arxiv.org/abs/2505.16068.

[24] Kaiyuan Zhang, Yuke Peng, Ke Jiang, and Yinqian Zhang. ActGov: Governing LLM agent actions via policy-constrained validation. arXiv preprint arXiv:2609.24446, 2026. URL https://arxiv.org/abs/2609.24446.