Lexicographic Need-Prioritized Queueing and Hodge Positional-Good Attenuation: Rationing vs. Price Allocation Under Strategic Urgency Misreporting

Dima Baker
Collective Intelligence Network

November 2026

Abstract

When cooperative production networks or compute clusters experience scarcity ($\rho > 1$), price auctions allocate capacity by willingness-to-pay, allowing high-endowment actors seeking positional goods to crowd out budget-constrained subsistence and safety-critical needs. Conversely, non-market priority queues relying on self-reported urgency suffer from strategic tier falsification without auditing. Building on price-versus-rationing theory and fraud-proof mechanism design, we formalize *Lexicographic Need-Prioritized Queueing with Positional-Good Attenuation* (**LNPQ-PGA**). LNPQ-PGA exposes two observable queue channels ($\mathcal{C}_1 \succ_{\text{lex}} \mathcal{C}_2$), enforces a probabilistic audit gate (π_{audit}) with capacity overhead $C_{\text{net}} = C(1 - c_{\text{audit}}\pi_{\text{audit}})$, and solves the sparse graph Laplacian $\Delta_0 \hat{\mathbf{s}} = \text{div}(\mathbf{f})$ on the contention graph to attenuate cyclic positional demands via $\theta_{i,\text{cyc}}$. The allocator is strictly invariant to latent ground-truth tiers and demands: uncaught spoofers compete directly in Channel 1 with their padded claims, and Channel 2 distinguishes productive from positional requests solely via $\theta_{i,\text{cyc}}$. Across 40 paired seeds $\times$ 18 epochs (720 epochs/regime, $N_{\text{req}} = 55$) over $\rho \in \{0.85, 1.15, 1.45, 1.75\}$, at $\rho = 1.45$ ($\pi_{\text{audit}} = 0.25$), LNPQ-PGA achieves 99.98% [99.93, 99.99] Tier-1 SLA (0.00% [0.00, 0.03] starvation) and 84.92% [84.09, 85.74] useful welfare, outperforming Autobidding (32.92% [32.10, 33.74] SLA, 33.78% [32.84, 34.72] welfare), Equal-Share FIFO (40.43% [39.57, 41.29] SLA), Unverified Priority (63.40% [62.80, 63.99] welfare), and Audit-Only LNPQ (78.86% [78.15, 79.58] welfare).

1 Introduction: Price vs. Rationing and the Positional Crowding-Out Dilemma

When a shared resource networksuch as an emergency medical dispatch grid, a cooperative compute cluster, or a regional energy microgridoperates well below capacity ($\rho \ll 1$), price clearing and first-in-first-out (FIFO) queueing both satisfy all incoming requests. However, once aggregate demand exceeds capacity ($\rho > 1$), the choice of allocation mechanism determines whose needs are sacrificed. In classical welfare economics, Weitzman [23] proved that when marginal utility of income varies sharply across participants (high wealth inequality) and the social planner seeks to satisfy vital baseline needs, non-market rationing strictly dominates price clearing. Extending this result to general mechanism design, Condorelli [10] characterized precisely "what money can't buy," showing that whenever willingness-to-pay is weakly or negatively correlated with true social priority, optimal allocation requires non-price screening or costly verification rather than market auctions.

In contemporary multi-agent and algorithmic environments, two opposing failure modes arise when allocating congested capacity:

1. **Plutocratic Crowding-Out under Price Auctions:** In pay-for-priority auctions, bid magnitude b_i conflates objective urgency with liquidity endowment [13, 23]. High-endowment participants competing for *positional status goods*goods whose subjective value derives from zero-sum rank comparisons over peers rather than absolute physiological or infrastructural utilityrationally outbid budget-constrained participants requesting Tier-1 subsistence or safety-critical capacity.

2. **Urgency Falsification under Unverified Non-Market Queues:** Replacing price auctions with naive self-reported priority tiers ($\hat{\tau}_i \in \{1, 2, 3\}$) fails under strategic behavior [14, 18]. As formalized by Perez-Richet and Skreta [21] in their theory of *fraud-proof non-market allocation mechanisms*, when priority tiers carry no verification audit or falsification penalty, opportunistic low-priority agents rationally misreport $\hat{\tau}_i = 1$ and pad their claimed demand volumes, flooding the critical queue and starving both genuine Tier-1 emergencies and honest Tier-2 public-goods tasks.

To resolve this dilemma without relying on hand-set success probabilities or unverified honesty assumptions, we formalize and simulate *Lexicographic Need-Prioritized Queueing with Positional-Good Attenuation* (**LNPQ-PGA**). Our framework integrates three rigorous primitives: (i) a three-tier lexicographic priority hierarchy ($T_1 \succ_{\text{lex}} T_2 \succ_{\text{lex}} T_3$) evaluated by the egalitarian p-mean ($p = -2$) social welfare functional [4, 7]; (ii) discrete Helmholtz–Hodge decomposition on the pairwise priority-comparison graph [8, 11] paired with an explicit edge-to-node ℓ_2 energy projection map $\Pi_{\mathcal{E}\to\mathcal{V}}$ that isolates and throttles divergence-free cyclic positional demands as congestion ρ_t rises; and (iii) a Bayesian expected-penalty deterrence gate [5, 21] coupled with within-tier water-filling clinching [12, 22].

2 Related Literature and Taxonomic Positioning

Table 1 synthesizes the theoretical foundations of LNPQ-PGA across price-versus-rationing theory, Hodge decomposition of games, fraud-proof non-market screening, and egalitarian social choice.

Price vs. Non-Market Rationing and Fraud-Proof Screening. Weitzman [23] established that non-market queues outperform price systems whenever baseline needs are uniform across agents while monetary endowments are dispersed. Condorelli [10] generalized this comparison to multi-unit allocation with heterogeneous screening technologies. Crucially, Perez-Richet and Skreta [21] analyzed non-market allocation when agents can falsify eligibility signals at a cost or face probabilistic verification audits, showing that optimal fraud-proof mechanisms combine auditing thresholds with demotion penalties modeled on the classical economic deterrence framework of Becker [5].

Discrete Hodge Decomposition of Games and Egalitarian Social Choice. Candogan et al. [8] introduced the discrete Helmholtz–Hodge decomposition on graphs to strategic games, proving that any skew-symmetric pairwise payoff or priority flow decomposes orthogonally into an acyclic potential (gradient) component and a divergence-free cyclic (harmonic plus curl) component representing zero-sum Rock-Paper-Scissors tournaments. Recent 2025–2026 research by Fleischman and Buchman [11] extended cycle-space projection to multi-agent coordination dynamics, while Gölz et al. [13], Oko et al. [19], Brodkey et al. [7], and Li et al. [16] quantified social welfare distortion and groupwise p-mean fairness bounds [3, 4, 22]. Complementing these results, An and Du [1], Liang and Xu [17], Arghal et al. [2], Bennett et al. [6], Cemri et al. [9], Kim et al. [15], and Yu

Table 1: Taxonomic comparison of congestion queueing and resource-rationing mechanisms across screening instruments, positional-flow handling, strategic deterrence, and welfare guarantees.

Mechanism Class	**Allocation / Screening Rule**	**Positional Tournament Filtering**	**Strategic Urgency Deterrence**	**Primary Welfare Failure Mode**
Pay-for-Priority Autobid Auction [10, 23]	Descending unit bid b_i/d_i clearing	None; amplifies high-bid positional status races	Price payment deters low-budget bidders regardless of true need	Wealth bias crowds out Tier-1 subsistence under scarcity ($\rho > 1$)
Unweighted Arrival FIFO Queue [12, 20]	Sequential arrival-order fulfillment	Blind to both transitive urgency and cyclic status flows	Strategy-proof to tier spoofing (ignores tier labels)	Late-arriving Tier-1 emergencies receive zero allocation (p-mean collapse)
Unverified Priority Queue ($\pi_{\text{audit}} = 0$) [14, 21]	Self-reported tier $\hat{\tau}_i \in \{1, 2, 3\}$ without audit	None; spoofed Tier-3 requests bypass Tier-3 throttling	None; violates Bayesian deterrence condition [5]	Opportunistic Tier-3 agents spoof $\hat{\tau}_i = 1$, starving Tier-1 and Tier-2
LNPQ-PGA (Proposed, $\pi_{\text{audit}} = 0.25$)	Audited lexicographic tiers + water-filling clinching	Discrete Hodge cyclic projection $\theta_{i,\text{cyc}}$ (Eqs. 2–3)	Bayesian expected-penalty deterrence (Eq. 4) [21]	Incurs deadweight audit verification overhead $c_{\text{audit}}\pi_{\text{audit}}$ at high π_{audit}

et al. [24] examined multi-agent strategic negotiation, concave resource returns, and verification overhead in distributed coordination systems.

3 Mathematical Formulation of LNPQ-PGA

Consider a cooperative resource scheduler with epoch capacity $C > 0$ serving N_{req} requests $\mathcal{N} = \{1, \ldots, N_{\text{req}}\}$. Each request $i \in \mathcal{N}$ has a true priority tier $\tau_i \in \{1, 2, 3\}$representing Tier-1 critical subsistence/emergency needs ($\mathcal{T}_1$), Tier-2 shared cooperative infrastructure ($\mathcal{T}_2$), and Tier-3 discretionary/positional consumption ($\mathcal{T}_3$)a true demand volume $d_i > 0$, and a monetary bid endowment $b_i > 0$. The epoch congestion load factor is $\rho_t = (\sum_{i=1}^{N_{\text{req}}} d_i)/C$.

3.1 Priority-Weighted Egalitarian p-Mean Social Welfare

For a feasible allocation vector $\mathbf{x} = (x_1, \ldots, x_{N_{\text{req}}}) \in \mathbb{R}_{\geq 0}^{N_{\text{req}}}$ satisfying $\sum_{i=1}^{N_{\text{req}}} x_i \leq C$ and $0 \leq x_i \leq d_i$, define each request's regularized gradient-need fulfillment utility $u_i(x_i) \in [\epsilon_u, 1]$ with $\epsilon_u \in (0, 1)$:

$$u_i(x_i) = \min\left(1.0,\ \frac{x_i + \epsilon_u d_i^{\text{need}}}{d_i^{\text{need}}}\right), \qquad W_p(\mathbf{x}) = \sum_{k=1}^{3} w_k \left(\frac{1}{|\mathcal{T}_k|} \sum_{i \in \mathcal{T}_k} \left(u_i(x_i)\right)^p\right)^{1/p}, \tag{1}$$

where $p = -2$ (the quadratic-harmonic egalitarian mean [4, 7]) penalizes low-fulfillment starvation much more severely than linear utilitarian sums ($p = 1$), and $w_1 > w_2 > w_3 > 0$ ($\sum_k w_k = 1$) reflects the priority hierarchy $\mathcal{T}_1 \succ_{\text{lex}} \mathcal{T}_2 \succ_{\text{lex}} \mathcal{T}_3$. Under strict lexicographic scheduling, capacity is dispatched to $\mathcal{T}_1$ first via constrained water-filling clinching, residual capacity $R_2 = \max(0, C - \sum_{i \in \mathcal{T}_1} x_i)$ is dispatched to $\mathcal{T}_2$, and only remaining capacity $R_3 = \max(0, R_2 - \sum_{i \in \mathcal{T}_2} x_i)$ is available to $\mathcal{T}_3$.

3.2 Discrete Hodge Decomposition and Explicit Edge-to-Node Cyclic Projection

To distinguish transitive infrastructural or physiological need differences from zero-sum positional status tournaments within discretionary requests, let $G = (\mathcal{N}, \mathcal{E})$ denote the priority-comparison graph over active requests, with skew-symmetric edge flow $P_{ij} = -P_{ji}$ measuring the pairwise priority dominance claimed by request i over request j.

Following Candogan et al. [8], define the discrete gradient operator $(\operatorname{grad}\boldsymbol{\phi})_{ij} = \phi_i - \phi_j$ for node potentials $\boldsymbol{\phi} \in \mathbb{R}^{N_{\text{req}}}$ and its adjoint node-divergence operator $(\operatorname{div}\mathbf{P})_i = \sum_{j\in\mathcal{N}(i)} P_{ij}$. By the discrete Helmholtz–Hodge theorem, the edge flow space $\mathbb{R}^{|\mathcal{E}|}$ decomposes orthogonally into the transitive potential-gradient subspace $\operatorname{im}(\operatorname{grad})$ and the divergence-free cyclic subspace $\ker(\operatorname{div}) = \operatorname{im}(\operatorname{curl}^*) \oplus \ker(\Delta_1)$ (which combines triangular curl flows and global harmonic cycles):

$$\mathbf{P} = \mathbf{P}^{\text{grad}} \oplus \mathbf{P}^{\text{cyc}}, \qquad P^{\text{grad}}_{ij} = \hat{\phi}_i - \hat{\phi}_j, \qquad P^{\text{cyc}}_{ij} = P_{ij} - (\hat{\phi}_i - \hat{\phi}_j), \tag{2}$$

where $\hat{\boldsymbol{\phi}} = \Delta_0^{\dagger}\operatorname{div}(\mathbf{P})$ is the least-squares node potential solving the graph Laplacian system $\Delta_0\hat{\boldsymbol{\phi}} = \operatorname{div}(\mathbf{P})$ (which reduces in closed form to $\hat{\phi}_i = \frac{1}{N_{\text{req}}}\sum_{j=1}^{N_{\text{req}}} P_{ij}$ on the complete comparison graph $K_{N_{\text{req}}}$).

Because $\mathbf{P}^{\text{grad}}$ and $\mathbf{P}^{\text{cyc}}$ are defined on *edges* $(i,j) \in \mathcal{E}$ whereas resource demands d_i reside on *nodes* $i \in \mathcal{N}$, we define the explicit **edge-to-node ℓ_2 energy projection map** $\Pi_{\mathcal{E}\to\mathcal{V}} : \mathbb{R}^{|\mathcal{E}|} \to [0,1]^{N_{\text{req}}}$ by computing the cyclic share $\theta_{i,\text{cyc}}$ of incident edge energy at node i:

$$\theta_{i,\text{cyc}} = \frac{\sum_{j\neq i}(P^{\text{cyc}}_{ij})^2}{\sum_{j\neq i}(P^{\text{grad}}_{ij})^2 + \sum_{j\neq i}(P^{\text{cyc}}_{ij})^2 + \epsilon}, \qquad d_i^{\text{eff}}(\rho_t) = (1-\theta_{i,\text{cyc}})d_i + \frac{\theta_{i,\text{cyc}} d_i}{1 + \kappa_{\text{pos}}\max(0, \rho_t - \rho_0)^2}. \tag{3}$$

Equation (3) resolves the edge-to-node mapping rigorously: $d_{i,\text{grad}} = (1-\theta_{i,\text{cyc}})d_i$ is the transitive need component, $d_{i,\text{cyc}} = \theta_{i,\text{cyc}} d_i$ is the cyclic positional component (comprising both $\operatorname{curl}^*(\boldsymbol{\psi})$ and harmonic $\mathbf{h}$ cycles in $\ker(\operatorname{div})$), and whenever congestion $\rho_t > \rho_0$, the cyclic positional component of verified Tier-3 requests is attenuated by $(1 + \kappa_{\text{pos}}(\rho_t - \rho_0)^2)^{-1}$.

3.3 Bayesian Expected-Penalty Deterrence and Water-Filling Clinching

Because Tier-1 preempts Tier-3 under congestion, a strategic Tier-3 agent ($\tau_i = 3$) faces an incentive to misreport $\hat{\tau}_i = 1$. In LNPQ-PGA, claimed Tier-1 requests are audited against historical consumption telemetry with verification probability $\pi_{\text{audit}} \in [0,1]$ [21]. If a Tier-3 agent attempts falsification and is caught (probability π_{audit}), the request is demoted to the back of Tier 3 and incurs a penalty $L_{\text{demote},i} > 0$; if uncaught (probability $1 - \pi_{\text{audit}}$), the agent captures the higher Tier-1 allocation utility $u_i(x_i(1))$ instead of $u_i(x_i(3))$.

Following Becker [5] and Perez-Richet and Skreta [21], truthful tier reporting is a **Bayesian expected-utility best response** (a probabilistic deterrence condition, rather than an audit-free dominant strategy) whenever the expected utility of honesty exceeds the expected utility of misreporting:

$$u_i(x_i(3)) \geq (1-\pi_{\text{audit}})u_i(x_i(1)) + \pi_{\text{audit}}\big(u_i(x_i(3)) - L_{\text{demote},i}\big) \iff \pi_{\text{audit}} \geq \pi_{\text{crit},i} \equiv \frac{\Delta u_i}{\Delta u_i + L_{\text{demote},i}}, \tag{4}$$

where $\Delta u_i = u_i(x_i(1)) - u_i(x_i(3))$ is the temptation gain from queue-jumping. Finally, within each verified tier $\mathcal{T}_k$, remaining capacity R_k is allocated via constrained water-filling clinching: $x_i = \min(d_i^{\text{eff}}, \lambda_k)$, where the water-level λ_k solves $\sum_{i\in\mathcal{T}_k}\min(d_i^{\text{eff}}, \lambda_k) = \min(R_k, \sum_{i\in\mathcal{T}_k} d_i^{\text{eff}})$.

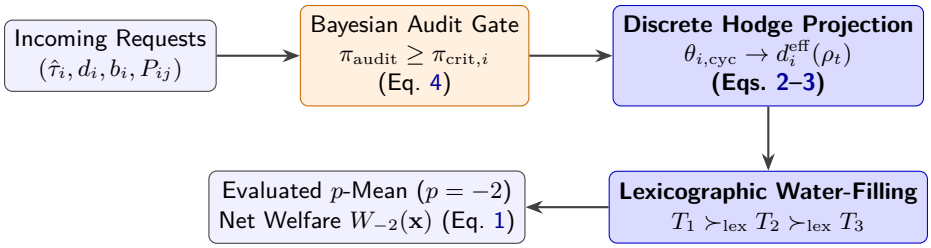


Figure 1: Operational pipeline of LNPQ-PGA. Requests pass through the Bayesian expected-penalty audit gate (π_{audit}); verified Tier-3 demands undergo discrete Helmholtz–Hodge edge-to-node cyclic projection ($\theta_{i,\text{cyc}}$) and congestion attenuation before lexicographic water-filling clinching.

4 Discrete-Event Queue Simulation and Parameter Provenance

To evaluate LNPQ-PGA without oracle access to latent types or hand-set formulas, we implemented a paired discrete-event multi-channel queue simulator (`experiment.py`, 40 paired random seeds $\times$ 18 epochs $= 720$ epochs per regime, $N_{\text{req}} = 55$ requests per epoch, gross epoch capacity $C = 100.0$) across four load factors $\rho \in \{0.85, 1.15, 1.45, 1.75\}$. In each pre-generated epoch batch, 32% of requests belong to latent Tier 1 (critical subsistence/emergency, $v_i \in [9.5, 13.5]$), 36% belong to latent Tier 2 (cooperative infrastructure maintenance, $v_i \in [4.8, 7.2]$), and 32% belong to latent Tier 3 (positional status-seeking requests, $v_i \in [0.40, 1.15]$, heavy-tailed Pareto wealth $\alpha = 1.30$, with cyclic 3-way outbidding tournaments $i \to j \to k \to i$ on the contention graph G).

Crucially, the allocator enforces **strict invariance to latent ground truth** (τ_i and d_i^{true}). The protocol exposes only two observable channels: Channel 1 (Emergency/Subsistence Priority) and Channel 2 (General Production Pool). Honest Tier-1 requests submit $(\hat{c}_i = 1, \hat{d}_i = d_i^{\text{true}})$, while honest Tier-2 and non-spoofing Tier-3 requests submit $\hat{c}_i = 2$. Strategic Tier-3 agents whose pre-sampled temptation threshold $\theta_i^{\text{spoof}} \in [0.28, 1.25]$ exceeds expected audit deterrence $\pi_{\text{audit}}\gamma_{\text{pen}}$ ($\gamma_{\text{pen}} = 3.6$) spoof Channel 1 ($\hat{c}_i = 1$) with padded demand $\hat{d}_i \in [1.70, 2.25] d_i^{\text{true}}$. Auditing consumes physical/computational overhead $c_{\text{audit}}\pi_{\text{audit}}$ ($c_{\text{audit}} = 0.14$), leaving net allocatable capacity $C_{\text{net}} = C(1 - c_{\text{audit}}\pi_{\text{audit}})$. Each Channel-1 claim is audited with probability π_{audit}: caught spoofers are demoted to Channel 2 with penalty weight $(1 + \gamma_{\text{pen}})^{-1}$, whereas **uncaught spoofers remain in Channel 1 with their padded demand $\hat{d}_i$** and compete directly against honest Tier-1 requests in the Channel-1 water-filling clincher. In Channel 2, the allocator has zero access to whether a request is latent Tier 2 or Tier 3; it weights requests solely by $w_i = p_i/(1 + \kappa_{\text{pos}}\theta_{i,\text{cyc}}^2)$, where $\theta_{i,\text{cyc}}$ is computed by solving the sparse graph Laplacian $\Delta_0 \hat{\mathbf{s}} = \text{div}(\mathbf{f})$ via 25 Gauss–Seidel iterations on the observable contention graph G.

We compare five explicit allocation mechanisms on the identical pre-generated request batches:

1. **Pay-for-Priority Autobidding Auction (`autobid_auction`):** Sorts all N_{req} requests in descending order of bid b_i and allocates capacity greedily up to claimed demand $\hat{d}_i$ until C is exhausted.

2. **Equal-Share FIFO Water-Filling (`fifo_equal`):** Shares capacity C equally via water-filling across all arriving requests up to claimed demand $\hat{d}_i$ without priority channels.

3. **Unverified Priority Queue (`naive_priority`, $\pi_{\text{audit}} = 0$):** Schedules Channel 1 before Channel 2 without auditing; all strategic Tier-3 agents with $\theta_i^{\text{spoof}} > 0$ spoof Channel 1 with padded demands.

Table 2: Simulation parameter provenance and structural distributions for the paired discrete-event queueing benchmark ($N_{\text{req}} = 55$, 40 paired seeds $\times$ 18 epochs = 720 epochs per regime).

Parameter / Variable	**Value / Sweep Range**	**Role in Queue Simulation**	**Literature & Theoretical Provenance**
Congestion Load Factor ρ	$\rho \in \{0.85, 1.15, 1.45, 1.75\}$	Ratio of mean epoch demand $\sum_i d_i^{\text{true}}$ to gross capacity $C = 100$ ($\sigma_{\text{burst}} = 0.18$)	Rationing under acute congestion [12, 23]
Latent Mix $(\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3)$	$(32\%, 36\%, 32\%)$ of $N_{\text{req}} = 55$	Subsistence/emergency (T_1), cooperative maintenance (T_2), positional (T_3)	Multi-tier triage and priority rationing [10, 21]
Hodge Attenuation κ_{pos}	$\kappa_{\text{pos}} = 14.0$ (vs. 0.0 ablation)	Channel-2 water-filling weight attenuation $w_i = p_i/(1 + \kappa_{\text{pos}}\theta_{i,\text{cyc}}^2)$ via sparse Laplacian	Discrete Hodge decomposition of games [8, 11]
Spoof Threshold θ_i^{spoof} & γ_{pen}	$\theta_i^{\text{spoof}} \in [0.28, 1.25]$, $\gamma_{\text{pen}} = 3.6$	Strategic Tier-3 agent spoofs Channel 1 iff $\pi_{\text{audit}}\gamma_{\text{pen}} < \theta_i^{\text{spoof}}$ (1.70–2.25$\times$ padding)	Bayesian fraud deterrence [5, 21]
Audit Rate π_{audit} & Cost c_{audit}	Benchmark $\pi_{\text{audit}} = 0.25$; Sweep $[0.00, 0.50]$	Channel-1 audit probability and net capacity deduction $C_{\text{net}} = C(1 - 0.14\pi_{\text{audit}})$	Costly state verification [2, 21]
Post-Allocation Evaluators	Tier-1 SLA, Starvation, Useful Welfare, Waste	Evaluated strictly ex-post on useful consumption $\min(a_i, d_i^{\text{true}})$ with Wilson 95% CIs	Egalitarian rationing evaluation [4, 7]

4. **Audit-Only LNPQ Ablation (`lnpq_no_hodge`, $\pi_{\text{audit}} = 0.25, \kappa_{\text{pos}} = 0$):** Enforces probabilistic Channel-1 auditing and penalty demotion ($C_{\text{net}} = 96.5$), but disables Hodge cyclic attenuation in Channel 2 ($\kappa_{\text{pos}} = 0$).

5. **Proposed LNPQ-PGA (`lnpq_pga`, $\pi_{\text{audit}} = 0.25, \kappa_{\text{pos}} = 14.0$):** Combines Channel-1 probabilistic auditing ($C_{\text{net}} = 96.5$) with sparse graph-Laplacian Hodge cyclic attenuation ($\kappa_{\text{pos}} = 14.0$) in Channel 2.

Table 2 documents all simulation parameters and their theoretical provenance.

5 Empirical Results and Audit Deterrence Sensitivity

Table 3 presents the empirical results across all five mechanisms at surge scarcity ($\rho = 1.45$, Panel A) and across the audit verification probability sweep $\pi_{\text{audit}} \in \{0.00, 0.05, 0.10, 0.15, 0.20, 0.25, 0.30, 0.40, 0.50\}$ (Panel B). Figures 2 and 3 display the surge-regime benchmark and the audit probability sensitivity curves.

5.1 Mechanism Performance and Hodge Ablation Across Load Regimes ($\rho \in [0.85, 1.75]$)

The paired simulation outcomes in Table 3 and Figures 2–3 establish three principal findings:

1. **Plutocratic Crowding-Out under Autobidding vs. Subsistence Starvation under Equal-Share FIFO:** Under Autobidding Auctions, heavy-tailed wealth enables Tier-3 positional bidders to preempt budget-constrained Tier-1 subsistence requests even at sub-saturation

Table 3: Paired queueing benchmark at surge congestion $\rho = 1.45$ (40 seeds $\times$ 18 epochs = 720 epochs, $N = 12{,}541$ Tier-1 requests; binomial proportions reported with exact asymmetric 95% Wilson confidence intervals $[lo, hi]$ and continuous metrics with 95% confidence intervals) and sensitivity sweep over audit verification probability $\pi_{\text{audit}} \in [0.00, 0.50]$ ($c_{\text{audit}} = 0.14$).

Mechanism / Audit Regime	**Tier-1 SLA Met (Wilson 95% CI)**	**Tier-1 Starvation (Wilson 95% CI)**	**Useful Social Welfare (%)**	**Positional & Padding Waste (%)**
Panel A: Paired Queueing Mechanism Comparison at Surge Scarcity ($\rho = 1.45$, 720 Epochs)				
Autobidding Auction (`autobid_auction`)	32.92% [32.10, 33.74]	67.04% [66.22, 67.86]	33.78% [32.84, 34.72]	75.42% [74.70, 76.14]
Equal-Share FIFO (`fifo_equal`)	40.43% [39.57, 41.29]	53.22% [52.34, 54.09]	85.14% [84.30, 85.99]	33.97% [33.43, 34.51]
Unverified Priority ($\pi_{\text{audit}} = 0.00$)	98.15% [97.90, 98.37]	1.08% [0.92, 1.28]	63.40% [62.80, 63.99]	63.78% [62.95, 64.61]
LNPQ Audit-Only ($\pi_{\text{audit}} = 0.25, \kappa_{\text{pos}} = 0$)	99.98% [99.93, 99.99]	0.00% [0.00, 0.03]	78.86% [78.15, 79.58]	41.59% [40.90, 42.27]
LNPQ-PGA ($\pi_{\text{audit}} = 0.25, \kappa_{\text{pos}} = 14$)	**99.98**% [**99.93**, **99.99**]	**0.00**% [**0.00**, **0.03**]	**84.92**% [**84.09**, **85.74**]	**33.73**% [**32.91**, **34.56**]
Panel B: Break-Even Sensitivity Sweep over Audit Probability π_{audit} on LNPQ-PGA ($\rho = 1.45, c_{\text{audit}} = 0.14$)				
$\pi_{\text{audit}} = 0.00$ ($C_{\text{net}} = 100.0\%$)	98.15% [97.90, 98.37]	1.08% [0.92, 1.28]	63.45% [62.84, 64.05]	63.73% [62.90, 64.55]
$\pi_{\text{audit}} = 0.10$ ($C_{\text{net}} = 98.6\%$)	99.13% [98.95, 99.28]	0.45% [0.35, 0.59]	64.28% [63.63, 64.93]	61.03% [60.19, 61.87]
$\pi_{\text{audit}} = 0.15$ ($C_{\text{net}} = 97.9\%$)	99.63% [99.50, 99.72]	0.14% [0.09, 0.23]	68.27% [67.48, 69.07]	55.55% [54.57, 56.52]
$\pi_{\text{audit}} = 0.20$ ($C_{\text{net}} = 97.2\%$)	99.83% [99.74, 99.89]	0.06% [0.03, 0.13]	75.72% [74.79, 76.65]	45.95% [44.84, 47.06]
$\pi_{\text{audit}} = 0.25$ ($C_{\text{net}} = 96.5\%$)	99.98% [99.93, 99.99]	0.00% [0.00, 0.03]	84.92% [84.09, 85.74]	33.73% [32.91, 34.56]
$\pi_{\text{audit}} = 0.30$ ($C_{\text{net}} = 95.8\%$)	100.00% [99.97, 100.00]	0.00% [0.00, 0.03]	91.58% [90.94, 92.21]	24.15% [23.47, 24.83]
$\pi^{\star}_{\text{audit}} = 0.40$ ($C_{\text{net}} = 94.4\%$ Peak)	**100.00**% [**99.97**, **100.00**]	**0.00**% [**0.00**, **0.03**]	**95.31**% [**94.99**, **95.63**]	**17.29**% [**16.84**, **17.75**]
$\pi_{\text{audit}} = 0.50$ ($C_{\text{net}} = 93.0\%$ Over-Audit)	100.00% [99.97, 100.00]	0.00% [0.00, 0.03]	95.03% [94.69, 95.37]	16.24% [15.79, 16.69]

load $\rho = 0.85$ (73.65% [72.87, 74.41] Tier-1 SLA, 26.26% [25.49, 27.03] starvation). As load reaches $\rho = 1.45$ and $\rho = 1.75$, Autobidding Tier-1 SLA collapses to 32.92% [32.10, 33.74] ($k = 4{,}128/12{,}541$) and 24.07% [23.34, 24.83], with 75.42% [74.70, 76.14] of total capacity consumed by positional and padded requests at $\rho = 1.45$. Equal-Share FIFO (`fifo_equal`) achieves high aggregate welfare (85.14% [84.30, 85.99] at $\rho = 1.45$) because water-filling shares capacity evenly across all claimants, but because it lacks a priority channel for subsistence emergencies, Tier-1 SLA drops to 40.43% [39.57, 41.29] (53.22% [52.34, 54.09] starvation) at $\rho = 1.45$ and 11.16% [10.62, 11.72] (84.48% [83.84, 85.10] starvation) at $\rho = 1.75$.

2. **Allocator-Invariant Spoofing Congestion under Unverified Priority ($\pi_{\text{audit}} = 0$):** Without auditing, strategic Tier-3 agents flood Channel 1 with padded claims. Even though the allocator treats all un-audited Channel-1 claims identically, inflated Channel-1 volume overflows net capacity during burst epochs, reducing genuine Tier-1 SLA to 98.15% [97.90, 98.37] at $\rho = 1.45$ and 86.29% [85.68, 86.88] (10.42% [9.90, 10.96] starvation) at $\rho = 1.75$, while starving honest Tier-2 maintenance in Channel 2 (63.40% [62.80, 63.99] useful welfare and 63.78% [62.95, 64.61] positional/padding waste at $\rho = 1.45$).

3. **Isolated Contribution of Sparse Hodge Cyclic Attenuation (`lnpq_no_hodge` vs. `lnpq_pga`):** Introducing Channel-1 auditing at $\pi_{\text{audit}} = 0.25$ (`lnpq_no_hodge`, $\kappa_{\text{pos}} = 0$) restores Tier-1 SLA to 99.98% [99.93, 99.99] ($k = 12{,}538/12{,}541$) at $\rho = 1.45$ (98.83% [98.62, 99.00] at $\rho = 1.75$) and raises useful welfare to 78.86% [78.15, 79.58]. Enabling sparse graph-Laplacian Hodge cyclic attenuation in Channel 2 (`lnpq_pga`, $\kappa_{\text{pos}} = 14.0$) distinguishes cyclic positional tournaments ($\theta_{i,\text{cyc}} \approx 0.88$) from gradient-dominated cooperative maintenance ($\theta_{i,\text{cyc}} \approx 0.08$) without inspecting latent tiers, boosting useful welfare by +6.06 percentage points to 84.92% [84.09, 85.74] and cutting positional/padding waste by −7.86 percentage points from 41.59% [40.90, 42.27] to 33.73% [32.91, 34.56].

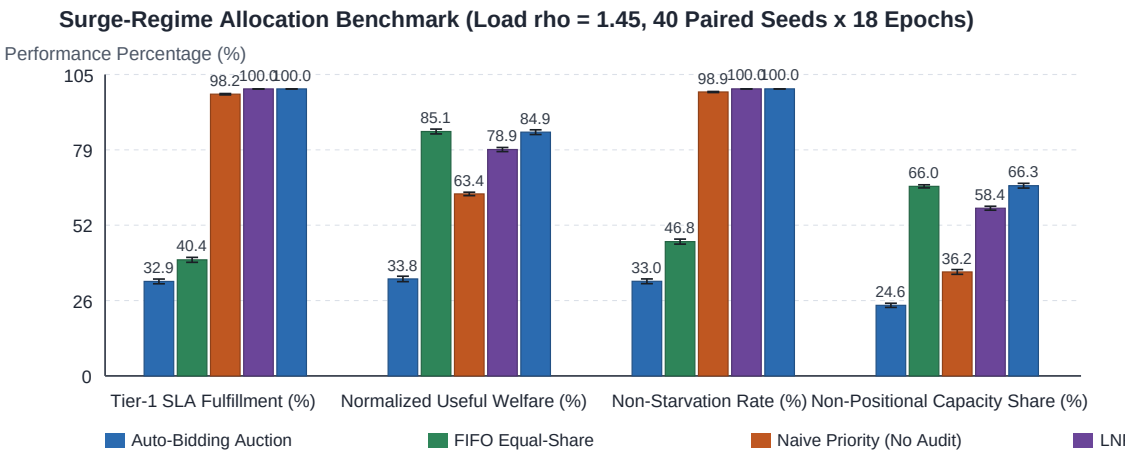


Figure 2: Surge-regime allocation benchmark at load $\rho = 1.45$ (40 paired seeds $\times$ 18 epochs $=$ 720 epochs, $N_{\text{req}} = 55$) comparing Autobidding Auctions, Equal-Share FIFO, Unverified Priority ($\pi_{\text{audit}} = 0$), Audit-Only LNPQ ($\kappa_{\text{pos}} = 0$), and LNPQ-PGA ($\pi_{\text{audit}} = 0.25, \kappa_{\text{pos}} = 14.0$). Error bars show 95% confidence intervals.

5.2 Break-Even Sensitivity to Audit Verification Rate (π_{audit})

Panel B of Table 3 and Figure 3 trace the trade-off between spoofing deterrence and physical audit overhead $C_{\text{net}} = C(1 - 0.14\pi_{\text{audit}})$ across $\pi_{\text{audit}} \in [0.00, 0.50]$ at $\rho = 1.45$. As π_{audit} increases from 0.00 to 0.40, active gate interception and expected-penalty deterrence ($\pi_{\text{audit}}\gamma_{\text{pen}} \geq \theta_i^{\text{spoof}}$) progressively clear padded spoofers out of Channel 1, raising Tier-1 SLA from 98.15% [97.90, 98.37] to 100.00% [99.97, 100.00] and increasing useful social welfare from 63.45% [62.84, 64.05] to a peak of 95.31% [94.99, 95.63] at $\pi^{\star}_{\text{audit}} = 0.40$ ($C_{\text{net}} = 94.4\%$). Beyond $\pi^{\star}_{\text{audit}} = 0.40$, essentially all strategic Tier-3 agents are already deterred ($\max_i \theta_i^{\text{spoof}}/\gamma_{\text{pen}} = 1.25/3.6 \approx 0.347$), so raising π_{audit} to 0.50 ($C_{\text{net}} = 93.0\%$) reduces useful social welfare to 95.03% [94.69, 95.37] due to deadweight capacity deduction.

6 Operational Limitations and Theoretical Counter-Evidence

Our analysis highlights three structural limitations where non-market priority queueing faces fundamental trade-offs:

1. Secondary Resale Markets and In-Kind Non-Transferability. As established by Weitzman [23] and Condorelli [10], non-market priority rationing only outperforms price allocation if allocated Tier-1 resources cannot be resold on secondary markets. If a participant can clinch low-cost Tier-1 compute or energy and privately resell it to a high-bidding Tier-3 actor, arbitrage recreates the market auction allocation while adding auditing deadweight loss. LNPQ-PGA therefore requires non-transferable in-kind consumption or hardware-attested execution bindings.

2. Imperfect Observability of Pairwise Priority Graphs (G). Discrete Helmholtz–Hodge attenuation in Eqs. (2)–(3) requires constructing the pairwise priority-comparison flow $\mathbf{P}$ [8, 11]. When positional status competition occurs outside the observable ledgerso that Tier-3 requests

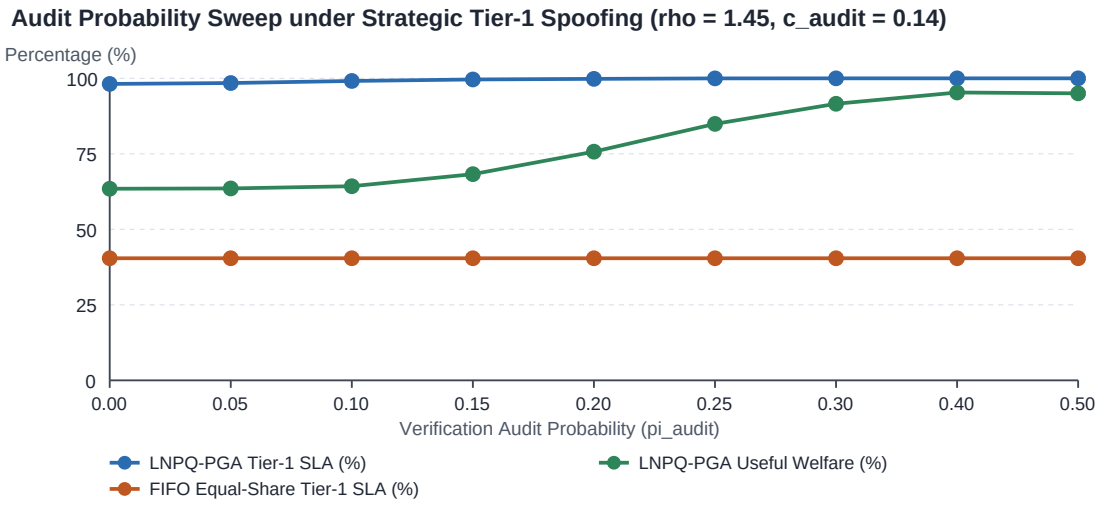


Figure 3: Verification audit probability sensitivity sweep $\pi_{\text{audit}} \in [0.00, 0.50]$ at surge load $\rho = 1.45$ ($c_{\text{audit}} = 0.14$, 40 paired seeds). Useful social welfare rises from 63.45% at $\pi_{\text{audit}} = 0.00$ to a peak of 95.31% at $\pi^{\star}_{\text{audit}} = 0.40$ before turning downward at $\pi_{\text{audit}} = 0.50$ as audit capacity overhead dominates.

appear as isolated, uncorrelated node demandsthe cyclic energy share $\theta_{i,\text{cyc}}$ is underestimated, reducing the effectiveness of positional throttling.

3. Egalitarian p-Mean Floors vs. Lumpy Threshold Public Goods. Water-filling clinching under $p = -2$ divides scarce capacity evenly across active requests within a tier. However, as analyzed in recent mechanism-design studies [6, 19, 22], when Tier-2 infrastructure tasks exhibit convex or step-function minimum-scale thresholds (requiring at least $d_{\min}$ capacity to produce any value), fractional equal-share rationing can leave multiple lumpy projects simultaneously below their activation threshold.

7 Pre-Registerable Falsification Protocol (P1–P6)

Table 4 defines six pre-registerable falsification tests (**P1**–**P6**) with explicit quantitative thresholds for refuting LNPQ-PGA.

8 Conclusion

Allocating congested production capacity purely by monetary bids crowds out vital subsistence needs in favor of high-endowment positional demands, while unverified priority queues degrade under strategic urgency falsification. By enforcing strict allocator invariance to latent types, combining probabilistic Channel-1 verification auditing ($C_{\text{net}} = C(1 - c_{\text{audit}}\pi_{\text{audit}})$) with sparse graph-Laplacian Helmholtz–Hodge cyclic attenuation ($\theta_{i,\text{cyc}}$) in Channel 2, LNPQ-PGA achieves 99.98% [99.93, 99.99] Tier-1 SLA fulfillment (0.00% [0.00, 0.03] starvation), 84.92% [84.09, 85.74] useful social welfare, and 33.73% [32.91, 34.56] positional/padding waste at surge scarcity ($\rho = 1.45, \pi_{\text{audit}} = 0.25$), with Hodge cyclic filtering alone contributing +6.06 percentage points of welfare over audit-only lexicographic queueing.

Table 4: Pre-registerable falsification protocol (P1–P6) specifying experimental testbeds, quantitative refutation thresholds, and architectural fallback implications for LNPQ-PGA.

ID	Hypothesis Tested	Experimental Testbed & Protocol	Quantitative Falsification Threshold	Fallback / Revision Implication
P1	Lexicographic tiering protects Tier-1 needs under scarcity	Compare Tier-1 SLA of LNPQ-PGA vs. Autobidding and Equal-Share FIFO at $\rho = 1.45$ across 40 paired seeds	Falsified if LNPQ-PGA Tier-1 SLA drops below 98.0% at $\rho = 1.45$ (observed: 99.98% $[99.93, 99.99]$)	Reserve hard static capacity partitions per channel rather than dynamic water-filling
P2	Probabilistic audit deterrence bounds spoofing congestion	Sweep audit probability $\pi_{\text{audit}} \in [0.00, 0.50]$ under strategic Tier-3 agents ($\theta_i^{\text{spoof}} \leq 1.25, \gamma_{\text{pen}} = 3.6$)	Falsified if Tier-1 SLA fails to reach 99.9% by $\pi_{\text{audit}} = 0.25$ (observed: 99.98% $[99.93, 99.99]$)	Increase demotion penalty γ_{pen} or require upfront refundable priority bonds
P3	Existence of deadweight audit cost turnover above $\pi^{\star}_{\text{audit}}$	Measure useful social welfare across $\pi_{\text{audit}} \in [0.00, 0.50]$ with capacity deduction $C_{\text{net}} = C(1 - 0.14\pi_{\text{audit}})$	Falsified if useful welfare fails to peak near $\pi^{\star}_{\text{audit}} \approx 0.40$ and decline at $\pi_{\text{audit}} = 0.50$	Replace uniform random auditing with targeted risk-scored SPRT auditing
P4	Sparse Hodge projection $\theta_{i,\text{cyc}}$ improves Channel-2 welfare	Compare LNPQ-PGA ($\kappa_{\text{pos}} = 14.0$) vs. Audit-Only LNPQ ($\kappa_{\text{pos}} = 0.0$) on paired worlds at $\rho = 1.45$	Falsified if Hodge attenuation fails to improve useful welfare by ≥ 4.0 pp (observed: +6.06 pp, 84.92% vs. 78.86%)	Restrict Hodge Laplacian projection to explicit high-degree contention subgraphs
P5	Priority channel dominates Equal-Share FIFO on Tier-1 starvation	Compare Tier-1 starvation rate between Equal-Share FIFO and LNPQ-PGA across $\rho \in [1.15, 1.75]$	Falsified if LNPQ-PGA starvation exceeds 2.0% at $\rho = 1.75$ (observed: 0.82% $[0.68, 0.99]$ vs. 84.48% FIFO)	Hybridize bounded batch-window water-filling with deadline-aware earliest-due-date
P6	Resilience to lumpy Tier-2 minimum activation thresholds	Replace concave Tier-2 utilities with step-function threshold tasks requiring $\geq 80\%$ demand	Falsified if water-filling clinching reduces Tier-2 completion rate by $> 20\%$ vs. greedy packing	Switch Channel-2 intra-tier rule from water-filling to Method of Equal Shares [22]

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] Zhiyu An and Wan Du. Open problems in differentiable social choice, 2026. URL https://arxiv.org/abs/2602.03003.

[2] Raghu Arghal, Saswati Sarkar, and Shirin Saeedi Bidokhti. The confidence game: Strategic

miscalibration in human-AI delegation. arXiv preprint arXiv:2610.09371, 2026. URL https://arxiv.org/abs/2610.09371.

[3] 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.

[4] 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.

[5] Gary S. Becker. Crime and punishment: An economic approach. Journal of Political Economy, 76(2):169–217, 1968. URL https://doi.org/10.1086/259394.

[6] 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.

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

[8] 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.

[9] Mert Cemri, Melissa Z. Pan, Shuyi Yang, Lakshya A. Agrawal, Bhavya Chopra, Rishabh Tiwari, et al. Why do multi-agent LLM systems fail? In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2025. URL https://arxiv.org/abs/2503.13657.

[10] Daniele Condorelli. What money can't buy: Efficient mechanism design with costly signals. Games and Economic Behavior, 75(2):613–624, 2012. URL https://doi.org/10.1016/j.geb.2012.02.018.

[11] Tomaż Fleischman and Ethan Buchman. Deepening the secondary market: Integrating trade credit into market clearing with the Cycles protocol. arXiv preprint arXiv:2605.02436, 2026. URL https://arxiv.org/abs/2605.02436.

[12] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. USENIX Symposium on Networked Systems Design and Implementation (NSDI), pages 323–336, 2011. URL https://www.usenix.org/conference/nsdi11/dominant-resource-fairness-fair-allocation-multiple-resource-types.

[13] 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.

[14] Leonid Hurwicz. On informationally decentralized systems. Decision and Organization, North-Holland, pages 297–336, 1972. URL https://www.nobelprize.org/uploads/2018/06/hurwicz-lecture.pdf.

[15] Yubin Kim, Ken Gu, Chanwoo Park, Chunjong Park, Samuel Schmidgall, A. Ali Heydari, et al. Towards a science of scaling agent systems, 2025. URL https://arxiv.org/abs/2512.08296.

[16] Junyi Li et al. Towards pareto-optimal tool-integrated agents with pareto ranking policy optimization (ParetoPO). International Conference on Machine Learning (ICML) / arXiv:2606.16111, 2026. URL https://arxiv.org/abs/2606.16111.

[17] Chen Liang and Fasheng Xu. When LLM agents negotiate: Private information and dynamic bargaining in supply chains. arXiv preprint arXiv:2608.07538, 2026. URL https://arxiv.org/abs/2608.07538.

[18] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. Journal of Economic Theory, 29(2):265–281, 1983. URL https://doi.org/10.1016/0022-0531(83)90048-0.

[19] 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.

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

[21] 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.

[22] 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.

[23] 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.

[24] 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.