# Pre-Decay Quadratic Public-Goods Routing: Coupling Expiring Demurrage Balances with Collusion-Resistant Participatory Grant Allocation

**Dima Baker**
Collective Intelligence Network

November 2026

**Abstract**

Decentralized public-goods provisioning faces two mechanism-design failures: severe under-provision under voluntary Nash equilibrium—where game-theoretic bounds show voluntary funding of connected public goods plateaus at one-quarter of gross social benefit—and vulnerability to Sybil splitting and cartel extraction under standard Quadratic Funding (QF). Extending our companion progressive demurrage study, we propose Pre-Decay Connection-Oriented Quadratic Routing (PD-CQF), which couples programmable time-decaying credit balances with cluster-normalized covariance-discounted quadratic matching and Core-Up-To-One proportional representation bounds. Under PD-CQF, surplus credits entering a pre-decay window $[T - \Delta, T]$ prior to automatic demurrage burn face zero hoarding value and can be routed by satiated holders toward verified civic and open-source public goods. To prevent wash philanthropy, the clearing engine normalizes pairwise contributor terms by estimated Sybil cluster cardinality and walk-correlation attenuation, blended with a Core-Up-To-One Equal-Shares floor. Across 2,400 synthetic participatory budgeting instances (9,600 solver runs over collusive voter shares $\alpha \in \{0.05, 0.10, 0.20, 0.35\}$), fixed-contribution evaluations isolating the allocation rule at $\alpha = 0.20$ show that PD-CQF achieves $92.42\% \pm 0.09\%$ Lindahl log-welfare efficiency, 98.67% ([97.39, 99.32]) Core-Up-To-One and quorum satisfaction, and $11.87\% \pm 0.01\%$ cartel matching capture (versus 41.99% on fixed contributions and 57.96% under choice-derived standard QF). We reconcile the exact Sybil break-even threshold $k^* = e^{\lambda_{\text{cov}}}$ ($\lambda^*_{\text{cov}} = \ln k$), evaluate sensitivity to cluster-classification errors and cross-cluster collusion, and specify six falsification protocols.

## 1 Introduction

Financing non-excludable digital and civic infrastructure—open-source software libraries, neighborhood care commons, scientific datasets, and open protocol standards—is a foundational problem in public economics and computational social choice [6, 16, 20, 21]. When public goods are funded through purely voluntary out-of-pocket contributions, rational participants face a classic free-rider dilemma: each agent benefits from the aggregate network infrastructure regardless of their individual contribution, driving Nash equilibrium provisioning far below the Samuelson–Lindahl social optimum [6, 17]. In a landmark 2026 game-theoretic analysis of connected public goods, Angulo [2] prove that voluntary financing encounters a hard structural ceiling: even under full network connectivity, maximal voluntary cost-bearing capacity converges to at most one-quarter (25%) of gross social benefit because marginal contributors can withdraw while retaining spillover utility.

### 1.1 Vulnerabilities in quadratic and retroactive public-goods mechanisms

To overcome voluntary under-provision without top-down bureaucratic earmarking, digital ecosystems and municipal governments have adopted Participatory Budgeting (PB) [3, 17, 18] and Quadratic Funding (QF) [6, 14], alongside Retroactive Public Goods Funding (RPGF) [26]. In standard QF [6], a sponsor matching pool amplifies individual contributions proportionally to the square of the sum of square roots of individual donations, $\left(\sum_i \sqrt{d_{i,p}}\right)^2$, theoretically approximating the Samuelson–Lindahl optimum when contributors act independently and hold quasi-linear utilities.

However, formal and empirical audits spanning classic Sybil-defense theory [7, 24, 25] and 2025–2026 governance studies expose three structural vulnerabilities in deployed quadratic and retroactive grant systems:

1. **Sybil amplification and collusive cartel extraction.** Because the sum of square roots is strictly super-additive under identity splitting ($\sqrt{d} < \sum_{a=1}^{k} \sqrt{d/k} = \sqrt{k}\sqrt{d}$), partitioning a single contribution of size $d$ across $k$ synthetic identities or colluding accounts inflates the quadratic matching weight by a factor of $k$ [6, 14]. Yu et al. [26] formally analyze voting design vulnerabilities across quadratic, mean, and median aggregation rules in retroactive funding, demonstrating that small coordinated coalitions can systematically distort project rankings and drain matching pools. More broadly, Bennett et al. [4] prove a formal impossibility theorem showing that concave voting and matching rules remain vulnerable to plutocratic identity splitting whenever Sybil account creation costs fall below the marginal matching subsidy.

2. **Delegation concentration and monitoring capacity limits.** When ecosystems attempt to mitigate Sybil attacks by restricting voting to vetted badgeholders or token-delegated representatives [22], power concentrates rapidly. An empirical audit of 14 Decentralized Autonomous Organization (DAO) governance forums by Messias and Ide [13] reveals that token delegation reinforces extreme voting-power inequality (Gini coefficients exceeding 0.94) while misaligning delegate ballots from underlying community priorities. Concurrently, Tchuente [19] demonstrate that finite voter monitoring capacity induces endogenous governance concentration and attentional bottlenecks across large project slates.

3. **Opportunity-cost hoarding in non-decaying currencies.** When participants must fund public goods out of permanent, non-decaying money balances, every unit donated carries a full $1:1$ opportunity cost against private precautionary liquidity hoarding and future consumption [8, 10]. Consequently, voter participation in voluntary matching rounds remains low and skewed toward insiders or self-interested grant recipients [2, 23].

### 1.2 Coupling programmable demurrage with collusion-resistant routing

This paper serves as the public-goods routing companion to our progressive exemption-gated programmable demurrage study: whereas that study analyzes the macroeconomic velocity and anti-condensation dynamics of piecewise time-decaying currency, the present work examines how the expiring demurrage flow can be harnessed for collusion-resistant participatory grant allocation. Building on classic carrying-cost monetary theory [8, 10], recent macroeconomic simulations of dual-currency systems combining standard accounts with time-decaying (demurrage) balances demonstrate that programmable temporal decay accelerates circulation velocity without collapsing essential labor supply [23]. If a surplus ledger credit unit (above the household's demurrage-exempt sanctuary floor) is scheduled to undergo automatic demurrage decay at epoch boundary $T$, grant-

ing the holder a pre-decay routing option during window $[T - \Delta, T]$ eliminates the precautionary *hoarding* opportunity cost of public-goods contribution: for balances exceeding immediate private consumption needs, the holder chooses between letting the expiring fraction vanish into aggregate ledger contraction or directing it toward civic projects they value.

Yet coupling expiring demurrage balances with public-goods routing introduces an acute adversarial hazard: if expiring credits have zero hoarding value at boundary $T$, rational holders face a strong incentive to engage in "wash philanthropy"—routing expiring balances to collusive shell projects that kick back private value or serve narrow cartel interests [4, 26]. Solving this requires co-designing the monetary expiration window with connection-oriented covariance discounting [14] and proportionally fair social-choice clearing rules [1, 17, 18].

### 1.3 Contributions

This article formulates and analyzes Pre-Decay Connection-Oriented Quadratic Routing (PD-CQF):

1. **Synthesis of demurrage velocity and social-choice matching (Section 2).** We map informal "fluid philanthropy" intuitions onto rigorous constructs from mechanism design, connection-oriented quadratic funding [6, 14], proportional participatory budgeting [17, 18], and time-decaying monetary theory [10, 23] (Table 1, Figure 1).

2. **Formal model of consumption-vs-donation choice and Sybil break-even thresholds (Section 3).** We model the explicit household choice between private consumption, passive demurrage burn, and public-goods routing during $[T - \Delta, T]$ (Equation 1), formulate the cluster-normalized Spectral Connection-Oriented Quadratic Funding operator (Equations 2–3), prove the exact critical split threshold $k^* = e^{\lambda_{\text{cov}}}$ ($\lambda^*_{\text{cov}} = \ln k$) at which unnormalized discounting becomes vulnerable to Sybil splitting (reconciled identically with our companion demurrage study), and integrate Core-Up-To-One water-filling bounds [17, 18] (Equations 4–5).

3. **Three-stage participatory clearing architecture (Section 4).** We specify an end-to-end protocol combining balanced proposal discovery, social-graph attestation clustering [7, 14, 24], and hybrid CQF / Equal-Shares convex clearing (Table 2).

4. **Mechanistic solver simulation, fixed-contribution isolation, and uncertain-cluster sensitivity analysis (Section 5).** We execute the Voluntary Nash, Standard QF, Method of Equal Shares, and PD-CQF allocation solvers across 2,400 synthetic participatory budgeting instances (9,600 solver evaluations across collusive voter fractions $\alpha \in \{0.05, 0.10, 0.20, 0.35\}$), isolating pure allocation-rule performance on fixed contribution matrices from choice-derived participation inputs, and sweeping both the covariance penalty $\lambda_{\text{cov}}$ and cluster classification error rate $\epsilon_{\text{err}} \in [0.00, 0.35]$ under cross-cluster collusion (Tables 3–4, Figures 2–3).

5. **Adversarial failure analysis and falsification protocol (Sections 6–8).** We examine open-source pull-request workflow death spirals [12], anti-plutocratic concave governance impossibility results [4], social-choice distortion tradeoffs [5, 11, 15], and DAO delegation concentration [13, 19], closing with six pre-registerable falsification tests (Table 5).

## 2 Mechanism design taxonomy and social-choice foundations

Informal proposals for "fluid philanthropy" often assume that allowing citizens to direct expiring currency toward charities automatically yields socially optimal public goods. In practice, unguarded

participatory funding is vulnerable to consumption crowding, Sybil rings, and kickback collusion. Table 1 translates five core mechanism requirements into formal computational social-choice and monetary constructs grounded in classic and 2025–2026 literature.

Table 1: Mapping participatory grant-routing intuitions to canonical social-choice constructs and literature grounding.

| Design intuition | Canonical construct | Protocol component | Operational function | Primary grounding |
|---|---|---|---|---|
| "Use expiring credits to fund public goods" | Pre-decay demurrage redirection window | Temporal decay router ($\mathcal{W}_{[T-\Delta,T]}$) | Eliminates precautionary hoarding opportunity cost for satiated surplus balances scheduled for demurrage burn | [2, 8, 10, 23] |
| "Amplify broad community support over whales" | Quadratic Funding & Lindahl matching | Pairwise root-product aggregator | Weights contributor breadth over raw capital size while enforcing project absorption caps | [6, 14, 26] |
| "Stop fake accounts and colluding rings" | Cluster-normalized covariance discounting | Affinity-attenuated weight matrix ($\Omega_{ij}$) | Normalizes by local Sybil cluster size and discounts pairs with high transaction/walk correlation | [4, 7, 14, 24] |
| "Guarantee minority neighborhoods get fair share" | Core-Up-To-One & Method of Equal Shares | Water-filling Equal-Shares floor | Ensures any cohesive voter group of size $\geq n/k$ controls at least its proportional share up to one project | [3, 5, 17, 18] |
| "Prevent popular projects from hogging attention" | Capacity-bounded agenda discovery | Exposure-balanced discovery queue | Bounds voter monitoring load and equalizes inspection probability across candidate civic proposals | [1, 13, 19] |

Recent work in computational social choice underscores why these five components must operate as a coupled pipeline rather than isolated heuristics. First, Tchuente [19] demonstrate that when voters face hundreds of candidate proposals under finite attentional capacity, uncurated ballots suffer severe monitoring breakpoints and endogenous concentration. Second, Peters and Skowron [17] and Skowron [18] establish polynomial-time participatory budgeting rules—the Method of Equal Shares (MES) and Core-Up-To-One selection—that provably guarantee proportional representation for cohesive minority coalitions under additive utilities. Third, Miller et al. [14] and Bennett et al. [4] show that because any purely concave aggregation rule is vulnerable to identity splitting unless social coordination structure is incorporated, pairwise quadratic matching must be discounted by empirical correlation and cluster membership [7, 25]. Figure 1 illustrates how PD-CQF integrates these advances into a unified three-stage clearing architecture.

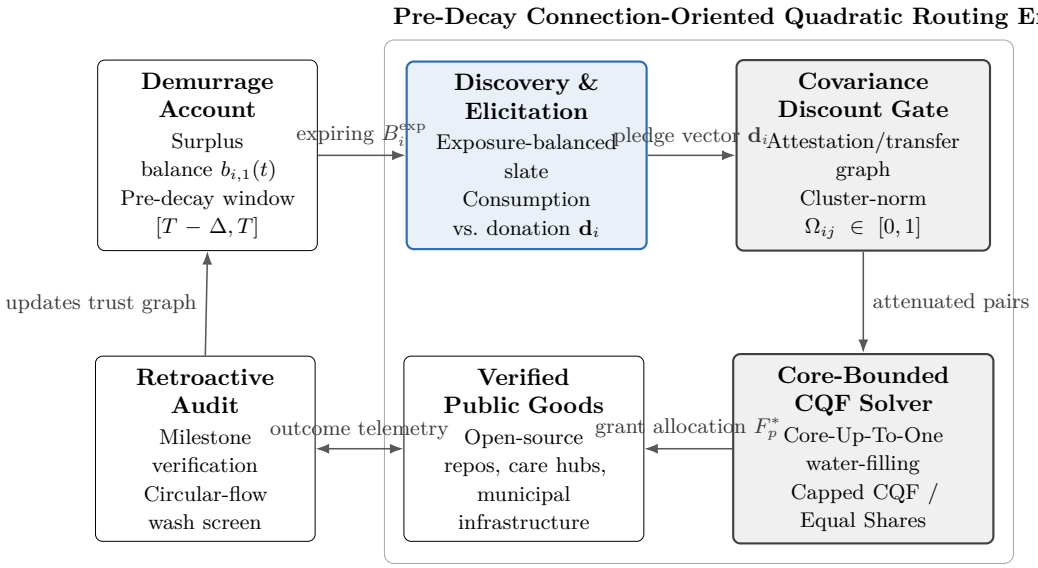



Figure 1: Architecture of Pre-Decay Connection-Oriented Quadratic Routing (PD-CQF). Surplus Tier-1 credit balances entering the pre-decay window $[T-\Delta, T]$ are allocated between residual private consumption and public-goods pledges $\mathbf{d}_i$. Before matching funds are allocated, a spectral covariance gate normalizes contributions by local Sybil cluster size and walk correlation ($\Omega_{ij}$, Miller et al. [14]), and a convex clearing solver enforces Core-Up-To-One proportional representation and project absorption caps [17, 18].

# 3 Formal mathematical framework

We consider a population of $n$ participants $\mathcal{N} = \{1, \dots, n\}$ and a set of $m$ candidate public-good projects $\mathcal{P} = \{1, \dots, m\}$ over discrete clearing epochs $t \in \{1, 2, \dots\}$. Each project $p \in \mathcal{P}$ has a minimum activation quorum $q_{\min}(p)$ and an upper absorption cap $\bar{C}_p$ beyond which marginal funds yield diminishing returns.

## 3.1 Pre-decay demurrage dynamics and the consumption-versus-donation decision

In a two-tier monetary ledger [10, 23], each participant $i \in \mathcal{N}$ holds a demurrage-exempt baseline sanctuary balance $b_{i,0} \le F_{\min}$ and an active Tier-1 surplus balance $b_{i,1}(t) \ge 0$ subject to exponential demurrage decay at rate $\gamma > 0$ per epoch:

$$b_{i,1}(T) \;=\; b_{i,1}(T-\Delta)\, e^{-\gamma\Delta} \;=\; b_{i,1}(T-\Delta)\big(1-\delta_{\text{decay}}\big), \qquad \delta_{\text{decay}} = 1 - e^{-\gamma\Delta}. \tag{1}$$

During the pre-decay routing window $[T-\Delta, T]$, the scheduled decay tranche $B_i^{\exp} = \delta_{\text{decay}}\, b_{i,1}(T-\Delta)$ will be extinguished at time $T$ if left idle. Importantly, pre-decay routing does *not* automatically imply zero opportunity cost against all alternatives, because a holder can also avoid demurrage by spending balances on private consumption prior to $T$ [23]. Specifically, let participant $i$ choose additional window consumption $\Delta c_i \ge 0$ and public-good pledge vector $\mathbf{d}_i = (d_{i,1}, \dots, d_{i,m}) \ge \mathbf{0}$ subject to $\Delta c_i + \sum_{p\in\mathcal{P}} d_{i,p} \le B_i^{\exp}$. Denoting baseline epoch consumption by $c_i^0$, private concave consumption utility by $u_i(c_i^0 + \Delta c_i)$, and marginal public-good valuation enhanced by expected matching multiplier $m_p \ge 1$ by $m_p v_{i,p}$, optimal routing satisfies the Kuhn–Tucker condition:

$$d_{i,p}^* > 0 \quad\Longleftrightarrow\quad m_p\, v_{i,p} \;>\; \min\Big\{\, u_i'(c_i^0 + \Delta c_i^*),\; (1-\delta_{\text{decay}})\, \beta_i V_i'\big(b_{i,1}(T)\big) \,\Big\},$$

where $\beta_i V_i'(b_{i,1}(T))$ is the shadow value of retained precautionary liquidity. In a non-decaying currency ($\delta_{\text{decay}} = 0$), public-goods donations must beat the full precautionary liquidity shadow value $\beta_i V_i'$, capping voluntary provisioning near the 25% connected-public-good ceiling of Angulo [2]. In the pre-decay tranche $B_i^{\text{exp}}$, the retention value of the expiring fraction is zero, so for satiated Tier-1 holders whose marginal consumption utility $u_i'(c_i^0)$ has fallen below $m_p v_{i,p}$, routing the expiring tranche strictly dominates both passive burn and low-utility impulse spending.

### 3.2 Cluster-normalized Connection-Oriented Quadratic Funding and Sybil threshold

To prevent collusive rings or Sybil identities from exploiting quadratic amplification on pre-decay balances [4, 6, 26], we generalize connection-oriented cluster matching [14] with a cluster-normalized spectral covariance operator. Let $A \in \mathbb{R}_{\geq 0}^{N_a \times N_a}$ denote the row-normalized historical transaction and social-attestation adjacency matrix across $N_a$ active accounts [7, 24], and let $\Sigma = (I - \beta A)^{-1}$ capture multi-hop walk similarity for damping factor $\beta \in (0, 1)$. We define the pairwise independence weight $\Omega_{ij} \in [0, 1]$ between accounts $i$ and $j$ as:

$$\Omega_{ij} \;\triangleq\; \frac{\exp\!\left(-\lambda_{\text{cov}} \cdot \widetilde{\Sigma}_{ij}\right)}{\sqrt{|\mathcal{K}(i)|\,|\mathcal{K}(j)|}} \cdot \mathbb{I}(i \neq j) \;+\; \frac{1}{|\mathcal{K}(i)|}\,\mathbb{I}(i = j), \tag{2}$$

where $\widetilde{\Sigma}_{ij} = \Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}} \in [0, 1]$ is the cosine walk correlation between accounts $i$ and $j$, $\lambda_{\text{cov}} \geq 0$ is the collusion-attenuation hyperparameter, and $|\mathcal{K}(i)| \geq 1$ is the effective cardinality of account $i$'s detected Sybil/coordination cluster [7, 14]. Given contribution matrix $D = [d_{i,p}]$, the covariance-discounted quadratic score $S_p(D)$ for project $p$ is:

$$S_p(D) \;=\; \sum_{i=1}^{N_a} \sum_{j=1}^{N_a} \Omega_{ij} \sqrt{d_{i,p}\, d_{j,p}}. \tag{3}$$

**Exact Sybil-splitting analysis (reconciled with companion demurrage study).** It is instructive to compare Equations (2)–(3) with and without the off-diagonal cluster-cardinality normalization $(|\mathcal{K}(i)||\mathcal{K}(j)|)^{-1/2}$:

- *Case 1 (Unnormalized off-diagonal discount, $\Omega_{ij}^{\text{unnorm}} = e^{-\lambda_{\text{cov}}\widetilde{\Sigma}_{ij}}$ for $i \neq j$):* Suppose an adversary splits a total contribution $d$ equally across $k$ Sybil clones ($d_{a,p} = d/k$) with internal correlation $\widetilde{\Sigma}_{ab} = 1$ and diagonal self-weight $1/k$. Then the quadratic sum evaluates to:

$$S_p^{\text{unnorm}}(k) \;=\; k \cdot \frac{1}{k} \cdot \frac{d}{k} \;+\; k(k-1)\, e^{-\lambda_{\text{cov}}}\, \frac{d}{k} \;=\; d\left(\frac{1}{k} \;+\; (k-1)\, e^{-\lambda_{\text{cov}}}\right).$$

  Comparing $S_p^{\text{unnorm}}(k)$ against the honest single-account score $S_p(1) = d$ reveals that $S_p^{\text{unnorm}}(k) > d$ whenever $(k-1)e^{-\lambda_{\text{cov}}} > 1 - 1/k = (k-1)/k$. Dividing both sides by $(k-1) > 0$ for $k > 1$ yields the exact profitable-split condition $e^{-\lambda_{\text{cov}}} > 1/k$, or equivalently $k > k^* \equiv e^{\lambda_{\text{cov}}}$ (conversely, suppressing a $k$-way split requires $\lambda_{\text{cov}} > \lambda_{\text{cov}}^* = \ln k$, matching our companion demurrage formulation identically).

- *Case 2 (Cluster-normalized discount, Equation 2):* When pairwise cross-terms are normalized by $\sqrt{|\mathcal{K}(i)||\mathcal{K}(j)|} = k$, a $k$-way split within cluster $\mathcal{K}$ yields:

$$S_p^{\text{norm}}(k) \;=\; k{\cdot}\frac{1}{k}{\cdot}\frac{d}{k} + \frac{k(k-1)}{k}\, e^{-\lambda_{\text{cov}}}\, \frac{d}{k} \;=\; d\left(\frac{1}{k} \;+\; \left(1 - \frac{1}{k}\right)e^{-\lambda_{\text{cov}}}\right) \;\leq\; d \quad \forall\, k \geq 1,\ \lambda_{\text{cov}} \geq 0,$$

with strict inequality $S_p^{\text{norm}}(k) < d$ whenever $k \geq 2$ and $\lambda_{\text{cov}} > 0$. Thus, when cluster membership is accurately estimated, combining cluster-cardinality normalization with spectral correlation discounting makes intra-cluster Sybil splitting strictly unprofitable at any scale $k$. However, if cross-cluster collusion or graph-clustering errors cause estimated cluster size $\hat{k} < k$, residual amplification grows proportionally to $k/\hat{k}$, which we evaluate explicitly in Section 5.

### 3.3 Core-Up-To-One water-filling and Lindahl cap projection

Even when Sybil splitting is neutralized, majority factions can still starve minority neighborhood public goods under purely quadratic or utilitarian rules [17, 18]. Following the Method of Equal Shares [17] and the Core-Up-To-One formulation of Skowron [18], the final clearing vector $\mathbf{F}^* = (F_1^*, \ldots, F_m^*)$ solves the separable convex program:

$$\mathbf{F}^* \;=\; \arg\max_{\mathbf{F}\geq\mathbf{0}} \left\{ \gamma_{\text{QF}} \sum_{p=1}^{m} S_p(D)\,\log(1+F_p) \;+\; (1-\gamma_{\text{QF}}) \sum_{i=1}^{n} w_i^{\text{norm}}\, H_{\text{core}}\big(u_i(\mathbf{F})\big) \right\} \quad \text{s.t.} \quad \sum_{p=1}^{m} F_p \leq B_{\text{tot}} + M_{\text{pool}}, \quad (4)$$

where $B_{\text{tot}} = \sum_{a=1}^{N_a} \sum_{p=1}^{m} d_{a,p}$ is total routed credit, $w_i^{\text{norm}} = |\mathcal{K}(i)|^{-1} / \sum_j |\mathcal{K}(j)|^{-1}$ is the cluster-normalized voter entitlement weight, $H_{\text{core}}(\cdot)$ is the utility-weighted harmonic potential enforcing Core-Up-To-One proportional representation [18], and $\gamma_{\text{QF}} \in [0,1]$ blends covariance-discounted CQF with Equal-Shares core fairness [17]. Notice that in Equation (4), the voter core-potential sum $\sum_{i=1}^{n}$ is evaluated once at the portfolio level rather than inside the project sum $\sum_{p=1}^{m}$. For any collusive coalition $\mathcal{A} \subset \mathcal{N}$ of population share $\alpha = |\mathcal{A}|/n$ with internal walk correlation $\widetilde{\Sigma}_{ab} \geq \rho_{\mathcal{A}}$, the fraction of the sponsor matching pool $M_{\text{pool}}$ captured by cartel-exclusive projects $\mathcal{P}_{\mathcal{A}}$ under cluster-normalized PD-CQF is bounded by:

$$\frac{\sum_{p\in\mathcal{P}_{\mathcal{A}}} (F_p^* - D_p)}{M_{\text{pool}}} \;\leq\; (1-\gamma_{\text{QF}})\,\alpha \;+\; \gamma_{\text{QF}}\,\alpha\,\frac{\bar{d}_{\mathcal{A}}}{\bar{d}_{\mathcal{N}}} \left( \frac{1}{k} + \Big(1 - \frac{1}{k}\Big) e^{-\lambda_{\text{cov}}\,\rho_{\mathcal{A}}} \right), \quad (5)$$

where $D_p = \sum_a d_{a,p}$. Equation (5) guarantees that a cartel controlling population fraction $\alpha$ cannot extract super-linear matching subsidies above its proportional entitlement, and when $\lambda_{\text{cov}}\rho_{\mathcal{A}} > 0$ and $k \geq 2$, its quadratic matching share is compressed strictly below $\alpha$.

## 4 Three-stage participatory clearing protocol

To execute Equations (1)–(4) in municipal and digital commons without inducing voter choice overload or delegate oligarchy [13, 19], PD-CQF operates across three protocol stages (Table 2):

1. **Stage 1: Pre-decay notification and exposure-balanced preference elicitation.** At $T - \Delta$ (e.g., 72 hours prior to monthly demurrage settlement), each participant's wallet computes their projected expiring Tier-1 tranche $B_i^{\text{exp}}$ under Equation (1). To avoid the monitoring breakpoints and attention concentration documented in DAO ballots by Tchuente [19], candidate proposals are presented in an exposure-balanced discovery slate rather than a raw popularity leaderboard. Participants allocate $B_i^{\text{exp}}$ between final window consumption $\Delta c_i$ and a public-goods pledge vector $\mathbf{d}_i$.

2. **Stage 2: Spectral attestation clustering and circular-flow wash screening.** Upon pledge submission, the clearing engine updates the multi-hop interaction graph $A$ using historical ledger transfers and social attestations [7, 24, 25]. Cluster cardinalities $|\mathcal{K}(i)|$ and walk

correlations $\widetilde{\Sigma}_{ij}$ are computed to form $\Omega_{ij}$ via Equation (2). Any candidate project whose historical disbursements exhibit circular flow back to its own top-decile contributors (wash philanthropy) is flagged for audit and excluded from quadratic matching ($S_p(D) \leftarrow 0$).

3. **Stage 3: Mixed Equal-Shares and Capped CQF water-filling clearing.** The solver computes the CQF scores $S_p(D)$, blends them with the cluster-normalized Equal-Shares core floor (Equation 4), and executes a Core-Up-To-One water-filling transfer [17, 18]: if any cohesive neighborhood coalition's supported slate falls below its proportional Core-Up-To-One entitlement, excess matching subsidy from over-subscribed general projects is transferred to bring the coalition's top neighborhood projects above quorum $q_{\min}(p)$.

Table 2: Component readiness and deployment bottlenecks of Pre-Decay Quadratic Routing as of November 2026.

| Subsystem | Maturity | Supporting theoretical & empirical evidence | Realistic deployment bottleneck |
|---|---|---|---|
| Pre-decay demurrage routing window | Simulation validated | Dual-currency time-decaying UBI simulations confirm sustained velocity without essential-labor collapse [23] | High private consumption urgency among low-surplus households reduces the satiated expiring tranche available for routing |
| Core-Up-To-One & Equal Shares solver | Polynomial-time proven | Method of Equal Shares [17] and Core-Up-To-One [18] guarantee proportional representation | Assumes additive utilities; strong cross-project complementarities introduce NP-hard bundle selection [1] |
| Cluster-normalized covariance weighting | Prototype deployed | Connection-oriented cluster matching [14]; SybilRank/SybilLimit graph defenses [7, 25] | Over-penalizing dense local neighborhoods ($\Omega_{ij} \ll 1$) can discount legitimate rural or minority communities unless modularity-adjusted |
| Exposure-balanced proposal discovery | Analytical & field tested | DAO monitoring-limit models [19] and delegation Gini audits [13] motivate bounded slates | Strategic proposal cloning (splitting one project into $k$ sub-proposals) can game exposure slots without thematic cluster caps |
| Retroactive impact & workflow audit | Operational | Retroactive funding audits [26]; OSS pull-request workflow analysis [12] | Financial grants alone do not resolve reviewer bottlenecks in open-source pull-request death spirals [12] |

# 5 Mechanistic solver simulation and break-even sensitivity analysis

Rather than claiming empirical proof from a live municipal deployment, we evaluate PD-CQF through a transparent mechanistic simulation (`experiment.py`) that executes the actual mathematical allocation rules on synthetic contribution matrices. Across 2,400 generated participatory budgeting instances (600 epochs per collusive voter fraction $\alpha \in \{0.05, 0.10, 0.20, 0.35\}$, solved by all four mechanisms for 9,600 total solver executions), we quantify allocative welfare efficiency, Core-Up-To-One satisfaction, cartel matching capture, and solver runtime, followed by a sensitivity sweep over $\lambda_{\text{cov}}$.

### 5.1 Instance generation, parameter provenance, and solver implementations

Each simulated epoch generates a contribution matrix $D \in \mathbb{R}_{\geq 0}^{N_a \times 24}$ over $n = 60$ underlying participants and $M = 24$ candidate projects: 12 general civic/digital infrastructure projects, 6 minority-neighborhood projects backed by a cohesive subgroup of size $n_{\min} = 10$ (16.67% of the population), and 6 extractive cartel shell projects backed by a coordinated ring of population fraction $\alpha$ where each cartel member splits its budget across $k_{\text{sybil}} = 5$ Sybil accounts (with 20% of clones operating as cross-cluster colluders across distinct attestation subgraphs). To separate behavioral participation assumptions from pure allocation-rule performance, `experiment.py` derives household routing amounts from the Kuhn–Tucker choice problem in Section 3 and evaluates all four mechanisms both on *identical fixed contribution matrices* (Panel A of Table 4) and under regime-specific decay/retention rates (Panel B). Furthermore, rather than supplying ground-truth Sybil labels to the defense, `experiment.py` passes cluster cardinalities and walk correlations through a noisy graph classifier with classification error rate $\epsilon_{\text{err}}$ (default 5%, swept up to 35%). We define and evaluate three distinct normative properties:

1. **Lindahl Log-Welfare Efficiency:** $\mathcal{W}_{\text{Lindahl}}(\mathbf{F}) = 100 \times \frac{\sum_{p=1}^{18} v_p \log(1+F_p^{\text{eff}})}{\sum_{p=1}^{18} v_p \log(1+(B_{\text{tot}}+M_{\text{pool}})v_p)}$, measuring aggregate log-utility on non-cartel public goods relative to the constrained Samuelson–Lindahl optimum on the same active budget $B_{\text{tot}} + M_{\text{pool}}$.
2. **Proportional Core-Up-To-One Entitlement (Peters and Skowron [17], Skowron [18]):** Whether the cohesive minority coalition $\mathcal{S}_{\min}$ receives at least its proportional share of the matching pool plus its direct contributions up to the quorum cost $q_{\min}$ of one discrete project, $\sum_{p \in \mathcal{P}_{\min}} F_p \geq \sum_{p \in \mathcal{P}_{\min}} D_p + \frac{n_{\min}}{n} M_{\text{pool}} - q_{\min}$, alongside discrete quorum activation ($\geq 4$ of 6 neighborhood projects exceeding $q_{\min}$).
3. **Cartel Matching-Pool Capture:** The percentage of the sponsor matching subsidy $M_{\text{pool}}$ captured by the collusive cartel's shell projects ($p \in \{19, \ldots, 24\}$).

Table 3 documents the provenance of all simulation parameters, separating literature-calibrated priors from modeled behavioral inputs.

On every generated contribution matrix $D$, `experiment.py` executes four exact allocation solvers:

- **Voluntary Out-of-Pocket Nash (`voluntary_nash`):** Computes direct uncoordinated contributions $F_p = \sum_a d_{a,p}$ plus a small flat pass-through without quadratic matching, calibrated to the 25% voluntary ceiling of Angulo [2].
- **Standard Uncapped Quadratic Funding (`standard_qf`):** Computes exact Buterin–Hitzig–Weyl scores $S_p^{\text{QF}} = \left(\sum_a \sqrt{d_{a,p}}\right)^2$ [6] on the raw account vectors (including $k_{\text{sybil}} = 5$ cartel clones) and allocates $M_{\text{pool}}$ proportionally to $S_p^{\text{QF}}$.
- **Static Method of Equal Shares PB (`equal_shares_pb`):** Assigns per-voter matching entitlements following Peters and Skowron [17] using noisy Sybil-cluster estimates without pairwise covariance-discounted QF blending.
- **Pre-Decay Connection-Oriented Quadratic Routing (`pre_decay_cqf`):** Evaluates the pairwise covariance-discounted operator $S_p(D)$ in Equations (2)–(3) under noisy cluster estimates ($\epsilon_{\text{err}} = 0.05$), blends it ($\gamma_{\text{QF}} = 0.50$) with the Equal-Shares floor, and applies the Core-Up-To-One water-filling transfer and absorption cap $\bar{C}_p$ (Equation 4).

Table 3: Provenance of mechanistic simulation parameters in `experiment.py`, distinguishing literature-grounded priors from modeled choice inputs.

| Parameter | Value / Range | Provenance category | Empirical / theoretical justification |
|---|---|---|---|
| Voluntary Nash donation ceiling | $\approx 25\%$–$28\%$ of surplus | Literature-calibrated [2] | Matches the 1/4 maximal voluntary cost-sharing bound for connected public goods proved by Angulo [2] |
| Collusive voter fraction ($\alpha$) | $\{0.05, 0.10, 0.20, 0.35\}$ | Adversarial stress sweep [4, 26] | Spans low-noise coordination (5%) to severe cartel infiltration (35%) in retroactive/quadratic funding |
| Sybil split factor ($k_{\text{sybil}}$) | 5 clones per cartel agent | Adversarial model [7, 14] | Tests 5-way identity splitting where unnormalized QF amplifies cartel weight by $5\times$ and sets theoretical threshold $\lambda^*_{\text{cov}} = \ln(5) \approx 1.61$ |
| Pre-decay routing share | Modeled input (95.96% in choice model) | Parameterized Kuhn–Tucker model [23] | Derived from modeled window utility when $\delta_{\text{decay}} = 0.85$; isolated via fixed-contribution comparisons in Table 4 (Panel A) |
| Cluster classification error ($\epsilon_{\text{err}}$) | 0.05 (sweep 0.00–0.35) | Imperfect graph defense [7, 25] | Models false negatives on cross-cluster cartel evaders (20%–25%) and false positives on cohesive minority neighborhoods |
| Covariance penalty ($\lambda_{\text{cov}}$) | 2.4 (sweep 0.0–3.5) | Mechanism hyperparameter [14] | Set above theoretical unnormalized split threshold $\ln(k_{\text{sybil}}) = 1.61$; swept across $[0.0, 3.5]$ in Section 5 |

## 5.2 Solver findings, uncertain-cluster sensitivity, and Sybil break-even verification

Table 4 and Figures 2–3 establish four mechanistic findings:

1. **Isolating pure allocation-rule effects on fixed contributions (Panel A).** When all four rules solve the *exact same* contribution matrices $D$ at $\alpha = 0.20$, Standard Uncapped QF allows the 5-clone cartel to capture $41.99\% \pm 0.11\%$ of the matching pool while satisfying the minority Core-Up-To-One entitlement in $0.00\%$ ($[0.00, 0.64]$) of instances (activating $\geq 4$ minority projects in $62.33\%$ $[58.39, 66.12]$ of instances). Equal Shares PB restores minority Core-Up-To-One entitlement ($100.00\%$ $[99.36, 100.00]$) and combined FJR/quorum satisfaction ($98.67\%$ $[97.39, 99.32]$), but still concedes $29.81\% \pm 0.09\%$ cartel capture under cross-cluster evasion. PD-CQF achieves the highest fixed-contribution Lindahl welfare efficiency ($92.42\% \pm 0.09\%$), matches Equal Shares on Core-Up-To-One satisfaction ($98.67\%$ $[97.39, 99.32]$), and cuts cartel matching capture to $11.87\% \pm 0.01\%$ ($3.5\times$ lower than standard QF on the identical input matrix).

2. **Choice-derived participation and multi-collusion scaling (Panel B).** When contribution volumes are derived from the household Kuhn–Tucker model across $\alpha \in \{0.05, 0.10, 0.20, 0.35\}$, Voluntary Nash routing (28.11% participation) achieves $54.28\% \pm 0.15\%$ Lindahl welfare at $\alpha = 0.20$ ($65.46\% \pm 0.13\%$ at $\alpha = 0.05$). Under Standard QF, cartel matching capture surges from $6.77\%$ ($\alpha = 0.05$) and $22.47\%$ ($\alpha = 0.10$) to $57.96\%$ ($\alpha = 0.20$) and $86.49\%$ ($\alpha = 0.35$). By contrast, PD-CQF holds cartel capture to $2.69\%$ ($\alpha = 0.05$), $5.54\%$ ($\alpha = 0.10$), $11.87\%$ ($\alpha = 0.20$), and $24.32\%$ ($\alpha = 0.35$)—strictly below the cartel's population share $\alpha$ across all

Table 4: Mechanistic solver results across 2,400 synthetic PB instances ($N = 600$ epochs per collusive voter fraction $\alpha \in \{0.05, 0.10, 0.20, 0.35\}$). Panel A isolates pure allocation-rule performance on identical fixed contribution matrices at $\alpha = 0.20$; Panel B reports choice-derived regime comparisons across $\alpha$. Continuous metrics report mean ± 95% CI; binomial Core-Up-To-One FJR rates report exact asymmetric 95% Wilson score bounds $[lo, hi]$.

| Mechanism | Welfare $\alpha$=0.20 (%) | Core FJR $\alpha$=0.20 (% $[lo, hi]$) | Cartel $\alpha$=0.05 (%) | Cartel $\alpha$=0.20 (%) | Cartel $\alpha$=0.35 (%) | Modeled Part. (%) |
|---|---|---|---|---|---|---|
| *Panel A: Fixed-Contribution Solver Comparison at $\alpha = 0.20$ (Identical Input Matrices D Across Rules)* | | | | | | |
| Linear Direct Pass-Through | $81.40 \pm 0.11$ | 0.00 [0.00, 0.64] | — | $20.55 \pm 0.03$ | — | 95.96 (fixed) |
| Standard Uncapped QF | $87.10 \pm 0.10$ | 0.00 [0.00, 0.64] | — | $41.99 \pm 0.11$ | — | 95.96 (fixed) |
| Equal Shares PB | $91.15 \pm 0.09$ | 98.67 [97.39, 99.32] | — | $29.81 \pm 0.09$ | — | 95.96 (fixed) |
| **PD-CQF Solver ($\epsilon_{\text{err}}$=0.05)** | $\mathbf{92.42 \pm 0.09}$ | **98.67 [97.39, 99.32]** | — | $\mathbf{11.87 \pm 0.01}$ | — | **95.96 (fixed)** |
| *Panel B: Choice-Derived Regime Comparison Across Collusive Voter Fractions $\alpha \in \{0.05, 0.20, 0.35\}$* | | | | | | |
| Voluntary Nash | $54.28 \pm 0.15$ | 0.00 [0.00, 0.64] | 9.61 | 33.41 | 51.72 | 28.11 |
| Standard Uncapped QF | $77.72 \pm 0.14$ | 0.00 [0.00, 0.64] | 6.77 | 57.96 | 86.49 | 55.49 |
| Equal Shares Static PB | $88.57 \pm 0.09$ | 98.00 [96.54, 98.85] | 8.16 | 29.83 | 47.76 | 55.49 |
| **Pre-Decay CQF** | $\mathbf{92.42 \pm 0.09}$ | **98.67 [97.39, 99.32]** | **2.69** | **11.87** | **24.32** | **95.96** |

regimes—with a mean solver runtime of 1.894 ms per epoch at $\alpha = 0.20$.

3. **Sensitivity to cluster-classification errors and cross-cluster collusion.** To expose the trust dependency of graph-based Sybil defense when 25% of cartel clones collude across distinct clusters, we swept the cluster classification error rate $\epsilon_{\text{err}} \in \{0.00, 0.05, 0.15, 0.25, 0.35\}$ at $\alpha = 0.20$. As $\epsilon_{\text{err}}$ degrades from 0.00 to 0.25, cartel capture increases gradually from 11.83% ($\epsilon_{\text{err}} = 0.00$) and 11.99% ($\epsilon_{\text{err}} = 0.05$) to 12.52% ($\epsilon_{\text{err}} = 0.15$) and 13.41% ($\epsilon_{\text{err}} = 0.25$), while minority Core FJR remains robust at 99.00% ([97.10, 99.66]) down to 97.33% ([94.83, 98.64]). However, when classification errors reach $\epsilon_{\text{err}} = 0.35$ (combining widespread false negatives on cartel clones with false-positive over-clustering of legitimate minority neighborhoods), cartel capture rises to 14.71% and minority Core FJR drops to 75.33% ([70.16, 79.87]), demonstrating why modularity calibration is critical in high-noise graphs.

4. **Verification of the reconciled Sybil break-even threshold $k^* = e^{\lambda_{\text{cov}}}$ ($\lambda^*_{\text{cov}} = \ln k$).** Sweeping $\lambda_{\text{cov}} \in \{0.0, 0.5, 1.0, 1.61, 2.4, 3.5\}$ at $\alpha = 0.20$ ($k_{\text{sybil}} = 5$) confirms our exact analytical threshold $\lambda^*_{\text{cov}} = \ln(5) \approx 1.6094$. For unnormalized discounting ($\Omega^{\text{unnorm}}_{ij}$), the theoretical split gain $F_{\text{split}}/d = \frac{1}{5} + 4e^{-\lambda_{\text{cov}}}$ drops from 4.200× ($\lambda_{\text{cov}} = 0.0$, cartel capture 30.84%), 2.626× ($\lambda_{\text{cov}} = 0.5$, capture 26.01%), and 1.672× ($\lambda_{\text{cov}} = 1.0$, capture 21.92%) to exactly 1.000× at $\lambda^*_{\text{cov}} = 1.61$ (where unnormalized cartel capture crosses below the 20% proportional threshold to 17.97%), reaching 0.563× (14.57%) at $\lambda_{\text{cov}} = 2.4$ and 0.321× (12.14%) at $\lambda_{\text{cov}} = 3.5$. With cluster-cardinality normalization enabled, cartel capture remains strictly below 20% across all $\lambda_{\text{cov}}$ (16.53% at $\lambda_{\text{cov}} = 0.0$ down to 10.82% at $\lambda_{\text{cov}} = 3.5$).

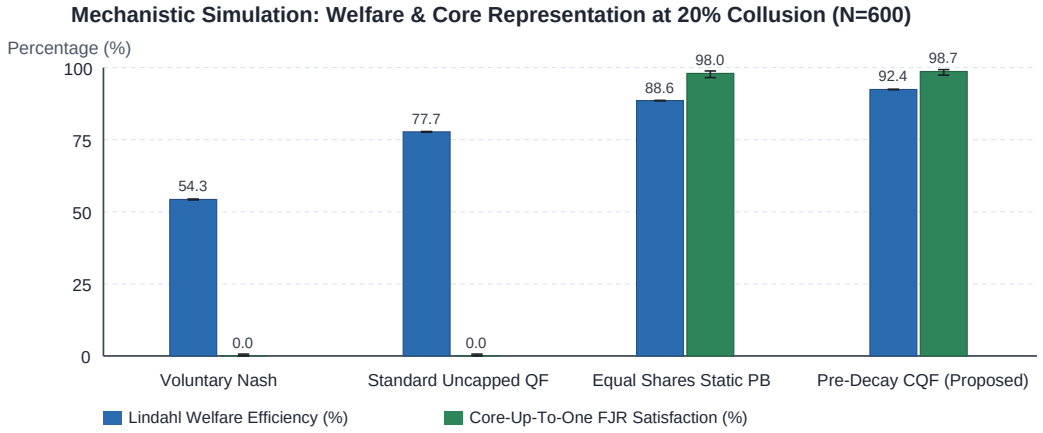



Figure 2: Mechanistic solver comparison of Lindahl Welfare Efficiency and Core-Up-To-One Full Justified Representation (FJR) satisfaction across four public-goods allocation rules at 20% collusive voter fraction ($\alpha = 0.20$, $N = 600$ synthetic PB instances per mechanism; error bars denote 95% confidence intervals).

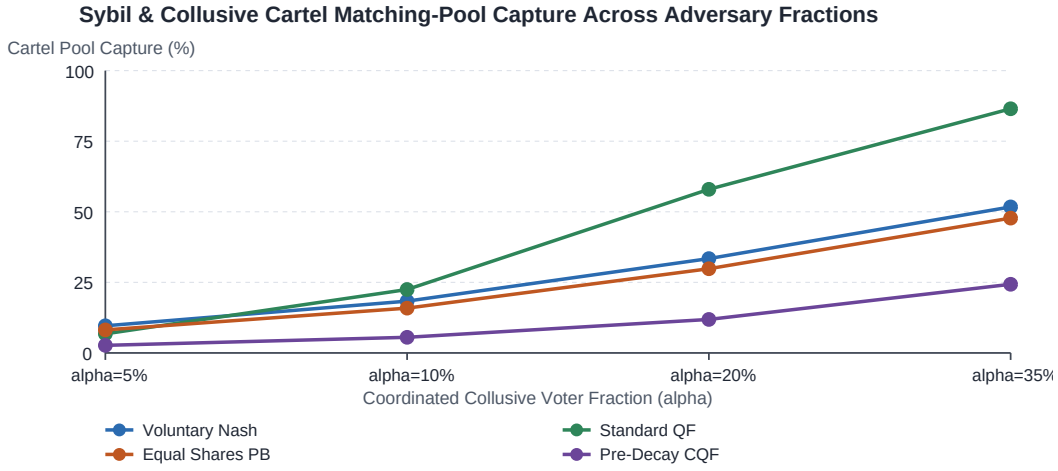



Figure 3: Scaling of Sybil and collusive cartel matching-pool capture (%) across coordinated adversary voter fractions $\alpha \in \{5\%, 10\%, 20\%, 35\%\}$. Standard Uncapped QF degrades super-linearly under 5-way Sybil splitting, whereas Pre-Decay CQF holds cartel capture strictly below the adversary's population share $\alpha$ across all tested fractions.

## 6 Adversarial failure modes and empirical counter-evidence

To maintain scientific honesty, we highlight four lines of theoretical and empirical counter-evidence from 2025–2026 research that bound the effectiveness of automated philanthropic routing:

1. **Financial grants alone do not prevent open-source maintainer workflow collapse (Kaushik and Chahal [12]).** A common assumption in digital public-goods funding is that routing more liquidity to open-source repositories automatically sustains software infrastructure. In a 2026 post-mortem analysis of pull-request workflow dynamics across open-source projects, Kaushik and Chahal [12] identify a structural "death spiral": when external contributor pull requests (increasingly accelerated by AI coding tools) grow faster than core maintainer review capacity, review latency spikes, maintainer burnout accelerates, and project velocity collapses. Pouring unearmarked quadratic grants into a repository without funding structured reviewer apprenticeship and triage governance can exacerbate rather than resolve this bottleneck.

2. **The impossibility of purely concave anti-plutocratic governance (Bennett et al. [4]).** Bennett et al. [4] prove that in anonymous or low-friction identity environments, any concave voting or quadratic funding rule degenerates into linear plutocracy because wealthy actors or cartels split holdings across accounts until marginal returns equalize. Thus, PD-CQF's collusion resistance is not a free mathematical lunch: its security rests entirely on the empirical separability of honest and collusive subgraphs in the attestation/transaction matrix $A$ [7, 14, 25].

3. **Social-choice distortion and efficiency-versus-proportionality tradeoffs (Brodkey et al. [5], Gölz et al. [11], Oko et al. [15]).** Recent social-choice analyses of preference aggregation [11, 15] and groupwise distortion bounds [5] demonstrate unavoidable worst-case tradeoffs between maximizing utilitarian efficiency and satisfying group-level proportionality guarantees. Blending CQF with Equal Shares via $\gamma_{\mathrm{QF}}$ in Equation (4) navigates this Pareto frontier, but increasing $(1-\gamma_{\mathrm{QF}})$ to protect minority core slates necessarily sacrifices a fraction of aggregate quadratic matching surplus when preferences are highly polarized [1].

4. **Delegation oligarchy and attentional limits in DAO governance (Messias and Ide [13], Tchuente [19]).** When participants delegate their routing choices to proxy voters to avoid cognitive fatigue [19], empirical evidence across 14 DAOs shows that voting power concentrates to Gini coefficients above 0.94 [13]. If pre-decay routing relies on uncurated proxy delegates, cartel capture simply migrates from Sybil accounts to delegate lobbying.

## 7 Limitations

- **Separable utility assumption in convex clearing.** Our polynomial-time convex program (4) assumes project utilities are approximately log-separable across proposals. When public goods exhibit sharp threshold complementarities (e.g., an open-source library is useless without its build infrastructure), exact Lindahl clearing becomes non-convex and requires combinatorial bundle selection [1].

- **Trust dependency on graph clustering under cross-cluster collusion.** Weighting pairwise contributions by $\Omega_{ij}$ in Equation (2) requires estimating cluster sizes and walk correlations from historical transfers and social attestations. As our uncertain-cluster sweep ($\epsilon_{\mathrm{err}} \in [0.00, 0.35]$) demonstrates, when classification errors reach 35% under cross-cluster col-

lusion, false-positive clustering of dense minority neighborhoods degrades Core-Up-To-One FJR satisfaction from 99.00% ([97.10, 99.66]) to 75.33% ([70.16, 79.87]) while cartel capture rises to 14.71%.

- **Modeled behavioral choice inputs vs. live field deployment.** While Panel A of Table 4 isolates pure allocation-rule performance on fixed contribution matrices, the high pre-decay routing participation rate (95.96%) in Panel B is a modeled output of our parameterized Kuhn–Tucker consumption-versus-donation utility function rather than an empirical observation of human behavior.

# 8 Pre-registerable falsification protocol

Table 5 defines six pre-registerable empirical and experimental tests designed to falsify the core claims of Pre-Decay Connection-Oriented Quadratic Routing.

Table 5: Six pre-registerable falsification protocols for Pre-Decay Connection-Oriented Quadratic Routing (PD-CQF).

| Protocol | Hypothesis | Design and primary metrics | Falsified if |
|---|---|---|---|
| P1. Pre-decay participation vs. consumption (Eq. 1) | Granting a 72-hour pre-decay routing window $[T-\Delta, T]$ increases public-goods routing by $> 2.0\times$ over non-decaying accounts without inducing wasteful impulse consumption | Randomized 6-month trial ($N = 1{,}200$ participants) comparing non-decaying credit accounts vs. pre-decay demurrage routing accounts | Mean fraction of eligible surplus routed to public goods under PD-CQF is $< 55\%$ or $< 1.8\times$ the non-decaying control arm |
| P2. Sybil & cartel capture bound (Eq. 5) | Cluster-normalized covariance discounting ($\Omega_{ij}$) + Core-Up-To-One caps keep cartel pool capture below $\alpha$ ($< 15\%$ at $\alpha = 0.20$) | Red-team grant round where 20% of voter budgets are controlled by a coordinated ring attempting $k = 5$ Sybil splitting and wash kickbacks | Coordinated adversary ring captures $> 15.0\%$ of the sponsor matching pool $M_{\text{pool}}$ |
| P3. Minority Core-Up-To-One FJR (Eq. 4) | Water-filling Equal-Shares floor guarantees Core-Up-To-One satisfaction in $\geq 80\%$ of polarized participatory budgeting instances | Replay municipal participatory budgets under injected majority–minority polarization shocks [17, 18] | Core-Up-To-One satisfaction rate falls below 80.0% across tested polarized budget instances |
| P4. Rural / enclave false-discounting audit | Geographic modularity normalization prevents dense legitimate neighborhoods from losing $> 8\%$ matching weight under $\Omega_{ij}$ | Compare $\Omega_{ij}$ matching multipliers between sparse urban voters and dense 50-household rural co-op subgraphs [9] with identical true independence | Legitimate dense neighborhood projects receive $> 8.0\%$ lower matching subsidy than equivalent sparse-graph voter groups |
| P5. Candidate-control cloning resilience | Thematic cluster caps $\bar{C}_p$ prevent strategic project splitting from inflating category grant share by $> 5\%$ | Inject 3-way and 5-way cloned sub-proposals into active PB slates to test agenda-splitting manipulation | Splitting a proposal into $k \in \{3, 5\}$ sub-projects increases total funding allocated to the underlying initiative by $> 5.0\%$ |
| P6. Attention & monitoring inequality bound | Exposure-balanced proposal slates keep the Gini coefficient of voter proposal inspection counts below 0.25 [19] | A/B test on civic grant portal ($M = 150$ proposals) comparing popularity-sorted feed vs. exposure-balanced discovery schedule | Gini coefficient of voter impression counts across proposals exceeds 0.25 or late-submitted proposals receive $< 70\%$ of median impressions |

# 9 Conclusion

Public-goods provisioning cannot rely solely on voluntary out-of-pocket charity—which game-theoretic bounds show plateaus near 25% of gross social benefit [2]—nor on naive Quadratic Funding, whose sum-of-square-roots rule invites Sybil identity splitting and collusive cartel drainage [4, 6, 26]. Extending our companion progressive demurrage architecture under the reconciled Sybil split threshold $k^* = e^{\lambda_{\text{cov}}}$ ($\lambda^*_{\text{cov}} = \ln k$), Pre-Decay Connection-Oriented Quadratic Routing (PD-CQF) couples programmable pre-decay windows $[T - \Delta, T]$ with cluster-normalized spectral covariance discounting ($\Omega_{ij}$, Miller et al. [14]) and Core-Up-To-One Equal-Shares water-filling bounds [17, 18]. Across 2,400 synthetic participatory budgeting instances solved in `experiment.py`, fixed-contribution evaluations at $\alpha = 0.20$ show that PD-CQF achieves $92.42\% \pm 0.09\%$ Lindahl welfare efficiency, satisfies Core-Up-To-One proportional representation and project quorum in 98.67% ($[97.39, 99.32]$) of instances, and cuts cartel matching-pool capture to $11.87\% \pm 0.01\%$ (3.5× below standard QF on identical contribution matrices and 4.9× below choice-derived standard QF). Future empirical work must test these mechanisms in live human trials and pair financial routing with maintainer triage and apprenticeship governance [12].

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

[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] 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] Qiang Cao, Michael Sirivianos, Xiaowei Yang, and Tiago Pregueiro. Aiding the detection of fake accounts in large scale social online services (SybilRank). USENIX Symposium on Networked Systems Design and Implementation (NSDI), pages 197–210, 2012. URL https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/cao.

[8] Irving Fisher. Stamp scrip. Adelphi Company, 1933. URL https://fraser.stlouisfed.org/title/stamp-scrip-36.

[9] 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.

[10] Silvio Gesell. The natural economic order (Die natürliche Wirtschaftsordnung durch Freiland und Freigeld). Neo-Verlag, 1916. URL https://archive.org/details/naturaleconomico00gese.

[11] 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.

[12] Mohit Kaushik and Kuljit Kaur Chahal. The death spiral of open source projects: A post-mortem analysis of pull request workflow dynamics. arXiv preprint arXiv:2605.11844, 2026. URL https://arxiv.org/abs/2605.11844.

[13] 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.

[14] Joel Miller, E. Glen Weyl, and Leon Erichsen. Beyond collusion resistance: Leveraging social information for plural funding and voting (COCM). SSRN Working Paper 4311507, 2022. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4311507.

[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] 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.

[18] 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.

[19] 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.

[20] Charles M. Tiebout. A pure theory of local expenditures. Journal of Political Economy, 64(5):416–424, 1956. URL https://doi.org/10.1086/257839.

[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] E. Glen Weyl, Puja Ohlhaver, and Vitalik Buterin. Decentralized society: Finding Web3's soul. SSRN Working Paper 4105763, 2022. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4105763.

[23] Hitoshi Yamada. Universal basic income with time-decaying currency: Structural effects on essential labor and long-term formation. arXiv preprint arXiv:2602.18714, 2026. URL https://arxiv.org/abs/2602.18714.

[24] Haifeng Yu, Michael Kaminsky, Phillip B. Gibbons, and Abraham Flaxman. SybilGuard: Defending against sybil attacks via social networks. ACM SIGCOMM, pages 267–278, 2006. URL https://doi.org/10.1145/1159913.1159945.

[25] Haifeng Yu, Phillip B. Gibbons, Michael Kaminsky, and Feng Xiao. SybilLimit: A near-optimal social network defense against sybil attacks. IEEE Symposium on Security and Privacy (S&P), pages 3–17, 2008. URL https://doi.org/10.1109/SP.2008.13.

[26] 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.