High-Dimensional In-Kind Clearing and Primal–Dual Convex Resource Allocation: Revisiting the Economic Calculation Debate Under Costly Verification

Dima Baker
Collective Intelligence Network

November 2026

Abstract

The economic calculation debate centers on a tension between incentive elicitation and liquidity constraints. Centralized convex planning and decentralized Lange–Arrow–Hurwicz primal–dual schemes solve the exact same underlying optimization problem under identical information, yet both suffer from strategic urgency inflation and endowment withholding when capacities are private. Spot fiat markets enforce hard budget constraints but freeze multi-hop supply chains when intermediate producers face working-capital bottlenecks. By the Hurwicz (1972) and Myerson–Satterthwaite (1983) impossibility theorems, no mechanism simultaneously achieves full allocative efficiency, incentive compatibility, individual rationality, and strict zero-cost budget balance. We formulate Audited Cycle-Compressed Primal–Dual Clearing (FP-PDIC), relaxing strict budget balance via costly physical stoichiometry audits alongside multilateral in-kind cycle cancellation. Stage 1 applies depth-first directed cycle cancellation (the Cycles Protocol) to settle reciprocal loops in physical units; Stage 2 resolves acyclic residuals via Augmented-Lagrangian ADMM; and Stage 3 applies randomized verification audits. Across paired $N = 40$-node convex economies and $D \in \{8, 16, 32, 64\}$ commodities (24 worlds per regime), unverified ADMM and centralized KKT programming converge to the identical distorted optimum (0.0003% gap; 92.34% [91.53, 93.14] efficiency at $D = 32$), while ADMM converges faster than Arrow–Hurwicz subgradients at 25 iterations (0.0007% vs. 0.5584% gap). Cash auctions reach 89.79% [89.29, 90.29] efficiency, and pure cycle barter degrades from 84.45% ($D = 8$) to 64.22% ($D = 64$). At audit rate $\kappa_{\text{audit}} = 0.20$, FP-PDIC reaches a break-even peak of 98.06% [97.94, 98.18] net efficiency before declining for $\kappa_{\text{audit}} \geq 0.30$.

1 Introduction

Whether a complex productive economy can coordinate high-dimensional resource allocation without relying on speculative fiat currency balances is one of the oldest foundational questions in mathematical economics and distributed systems. In the classical calculation debate, Lange [17] and Arrow and Hurwicz [3] demonstrated that a central planning board—or a decentralized network iterating over trial-and-error shadow prices (Lagrange multipliers)—can replicate the static first-order optimality conditions of a competitive general equilibrium whenever production sets are convex and local utilities are concave. In sharp response, Hayek [13] argued that the core economic problem is not the computational resolution of a known system of equations, but the dispersed, tacit nature of local knowledge regarding shifting capacities, technological stoichiometry, and urgent needs. Formalizing Hayek's critique, Hurwicz [14] and Myerson and Satterthwaite [20] proved that

in environments with privately held valuations and endowments, no informationally decentralized mechanism can simultaneously satisfy Pareto allocative efficiency, Bayesian or dominant-strategy incentive compatibility (IC), individual rationality (IR), and strict budget balance.

Yet purely market-mediated spot coordination suffers from a dual structural pathology when applied to multi-hop industrial and computational supply chains. In a spot double auction mediated by scalar currency, an intermediate node can only acquire upstream inputs if it currently holds liquid working capital. As Fleischman and Buchman [10] demonstrate with the Cycles Protocol, networks of productive enterprises routinely experience liquidity gridlock: every firm in a closed directed loop possesses idle physical capacity demanded by its downstream neighbor, yet multi-hop trade freezes because intermediate participants lack pre-existing cash balances. At the same time, empirical studies of autonomous language-model and algorithmic marketplaces in 2025–2026 show that unmediated negotiation under private information induces strategic delay, miscalibration, topology coordination failures, and epistemic herding [2, 8, 9, 16, 18, 19].

1.1 Separating solver equivalence from incentive and liquidity regimes

A rigorous modern treatment of non-fiat resource clearing must begin by avoiding a common conceptual confusion between *optimization algorithms* and *information/incentive regimes*:

1. **Mathematical equivalence of Centralized Convex Programming and Decentralized ADMM.** For any fixed vector of reported local urgency parameters $\hat{\boldsymbol{\theta}}_i \in \mathbb{R}_+^D$ and reported physical endowments $\hat{\mathbf{s}}_i \in \mathbb{R}_+^D$, a centralized convex planner solving the global resource allocation program via Karush–Kuhn–Tucker (KKT) water-filling and a decentralized network solving the same program via Alternating Direction Method of Multipliers (ADMM) or Arrow–Hurwicz dual gradient updates [3, 22] converge to the *exact same optimal allocation* $\mathbf{x}^\star(\hat{\boldsymbol{\theta}}, \hat{\mathbf{s}})$. Decentralizing the solver reduces central communication bottlenecks, but it cannot by itself outperform the centralized program it solves.

2. **The Hurwicz–Myerson–Satterthwaite impossibility trade-off.** If local parameters $(\boldsymbol{\theta}_i, \mathbf{s}_i)$ are private and unverified, rational best-response agents in a non-market allocation system have a strict incentive to exaggerate input urgency ($\hat{\theta}_{i,d} > \theta_{i,d}$) and withhold local endowment ($\hat{s}_{i,d} < s_{i,d}$) for private buffer consumption [14]. Both centralized planning and unverified ADMM converge to the *same distorted equilibrium*. To restore truthfulness without reintroducing speculative cash constraints, the mechanism must explicitly relax strict zero-cost budget balance by expending a fraction of social surplus on costly verification audits, as formalized in the fraud-proof non-market allocation theory of Perez-Richet and Skreta [24] and the economic deterrence framework of Becker [5].

3. **Topological sparsity of pure in-kind barter cycles in high dimensions.** Multilateral graph-cycle clearing [10] cancels closed loops of reciprocal obligations directly in kind without requiring shadow prices or cash liquidity. However, as commodity heterogeneity scales from $D = 8$ to $D = 64$ distinct resource classes, exact stoichiometric coincidence around closed cycles grows sparse, stranding acyclic net imbalances unless paired with a primal–dual shadow clearing stage [4, 11].

1.2 Contributions

In this paper, we formulate and evaluate Audited Cycle-Compressed Primal–Dual Clearing (FPPDIC), an analytical and simulation study of hybrid non-fiat resource coordination:

1. **Calculation Debate Synthesis and Impossibility Mapping (Section 2).** We map the classical Lange–Hayek–Hurwicz debate onto modern mechanism-design and distributed-systems primitives, stating explicitly which axiom of the Hurwicz/Myerson–Satterthwaite triad is relaxed by each institutional regime (Table 1).

2. **Coupled Cycle-Cancellation and Audited Primal–Dual Formulation (Section 3).** We formalize resource clearing over an N-node, D-resource production network. Stage 1 removes divergence-free in-kind trade cycles $\mathbf{F}^{\star}_{\text{cyc}}$ [10]; Stage 2 resolves residual imbalances via concave KKT/ADMM shadow multipliers $\boldsymbol{\lambda} \in \mathbb{R}^{D}_{+}$ [3, 11]; and Stage 3 enforces a Bayesian fraud-proof deterrence condition via costly sequential verification audits [5, 24, 26] (Equations 2–5).

3. **Protocol Architecture and Parameter Provenance (Section 4).** We specify the three-stage clearing architecture and document the theoretical and empirical provenance of all model parameters (Figure 1, Table 2).

4. **Solved Convex Allocation, Solver Convergence, and Break-Even Audit Sensitivity Benchmark (Section 5).** In `experiment.py`, we construct and solve actual paired 40-node, D-resource concave allocation instances ($D \in \{8, 16, 32, 64\}$, 24 paired worlds per dimension) via exact KKT bisection, true Augmented-Lagrangian ADMM, and Arrow–Hurwicz subgradient ascent under best-response strategic misreporting, comparing five regimes and sweeping the verification audit rate $\kappa_{\text{audit}} \in [0.00, 0.45]$ to identify the exact net-welfare break-even threshold (Table 3, Figures 2–3).

5. **Adversarial Failure Modes and Pre-Registerable Falsification Protocol (Sections 6–8).** We examine non-convex indivisibilities, delayed verification instability [15], and collusion risks, and provide six falsifiable test protocols (Table 4).

2 Institutional taxonomy and mechanism-design trade-offs

Discussions of non-fiat resource allocation frequently conflate three distinct functions: (i) the *medium of settlement* (whether reciprocal obligations require pre-existing liquid currency to clear), (ii) the *computational solver* (centralized batch optimization versus decentralized primal–dual iteration), and (iii) the *incentive verification regime* (whether private urgency and endowment reports are disciplined by spot budget constraints or costly physical audits). Table 1 situates five canonical coordination primitives within foundational pre-2025 economic theory and 2025–2026 multi-agent systems literature.

As summarized in Table 1, the fundamental difference between a centralized convex planner and FP-PDIC is *not* that ADMM magically discovers a better optimum than the linear or convex program it decomposes. Rather, an unverified planner (whether solved centrally via interior-point/KKT methods or decentrally via ADMM) violates incentive compatibility under Hurwicz [14], inducing rational participants to overstate need $\hat{\theta}_{i,d}$ and hide supply $\hat{s}_{i,d}$. Conversely, a spot market enforces budget balance and truthfulness via hard cash-in-advance constraints, but sacrifices allocative efficiency whenever credit or working-capital limits bind along multi-hop supply chains [10, 18]. FP-PDIC explicitly occupies the costly-verification regime of Perez-Richet and Skreta [24]: it eliminates cash-in-advance deadweight loss via in-kind cycle compression and shadow clearing, while paying an explicit physical audit cost $c_{\text{audit}}\kappa_{\text{audit}}$ to bound strategic falsification. Figure 1 illustrates this three-stage pipeline.

Table 1: Mapping resource-coordination mechanisms across the calculation debate and Hurwicz–Myerson–Satterthwaite impossibility trade-offs.

Design dimension	Classical foundation	FP-PDIC primitive	Impossibility trade-off & operational role	Primary grounding
Multilateral in-kind clearing	Clearinghouse netting & graph circulation	Cycle cancellation operator ($\mathcal{P}_{\text{cyc}}$)	Cancels closed reciprocal obligation loops in physical units without working-capital cash constraints	[4, 10]
Shadow scarcity valuation	Lange–Arrow–Hurwicz trial-and-error duals	Ephemeral KKT/ADMM duals ($\boldsymbol{\lambda} \in \mathbb{R}_+^D$)	Converges to exact centralized convex optimum under identical reports; non-hoardable across epochs	[3, 11, 17, 22]
Strategic report deterrence	Hurwicz & Myerson–Satterthwaite impossibility	Fraud-proof costly audit gate ($\kappa_{\text{audit}}, \Phi_i$)	Relaxes zero-cost budget balance: expends audit cost $c_{\text{audit}}\kappa_{\text{audit}}$ to deter urgency/endowment fraud	[5, 13, 14, 20, 24]
Subsistence & multi-axis fairness	Nash bargaining & proportional fairness	Priority-weighted concave utility ($U_i(\mathbf{x}_i)$)	Protects critical infrastructure and subsistence nodes against working-capital exclusion	[7, 12, 21, 25]
Sequential telemetry verification	Sequential probability ratio testing (SPRT)	Cross-signed receipt martingale ($E_{i,t}$)	Detects systematic input-output stoichiometry drift across multi-hop supply chains	[15, 19, 23, 26]

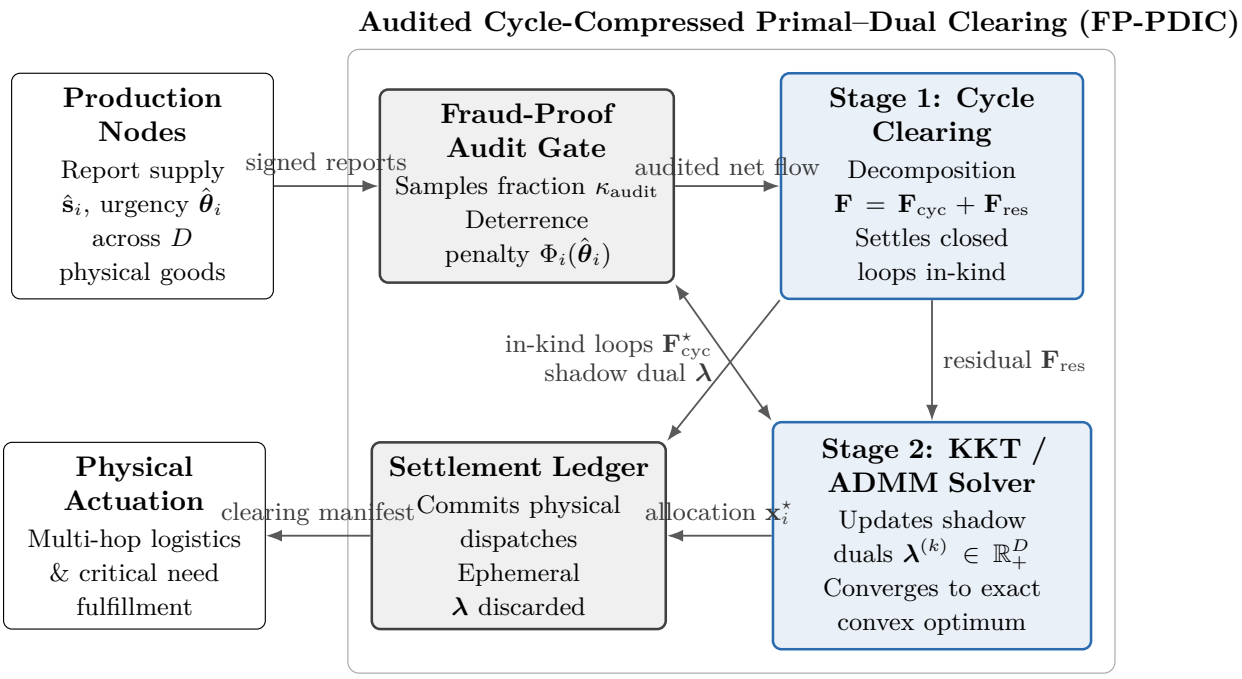


Figure 1: Architecture of Audited Cycle-Compressed Primal–Dual Clearing (FP-PDIC). Cooperative production nodes submit multi-dimensional urgency $\hat{\boldsymbol{\theta}}_i$ and supply $\hat{\mathbf{s}}_i$ reports across D commodities. A fraud-proof audit gate inspects a fraction κ_{audit} of physical conversion receipts before Stage 1 cancels closed reciprocal obligation loops in kind via the Cycles Protocol. Stage 2 resolves residual net imbalances using decentralized ADMM shadow duals $\boldsymbol{\lambda} \in \mathbb{R}^D_+$ that converge to the exact KKT optimum of the reported convex program and vanish upon epoch settlement.

3 Formal mathematical framework

We consider a productive network of N cooperative nodes $\mathcal{V} = \{1, \dots, N\}$ interacting across D distinct physical and computational commodities $\mathcal{D} = \{1, \dots, D\}$. At clearing epoch t, each node $i \in \mathcal{V}$ possesses a true physical endowment vector $\mathbf{s}_i = (s_{i,1}, \dots, s_{i,D})^\top \in \mathbb{R}^D_+$ and a true concave production/consumption utility function:

$$U_i(\mathbf{x}_i; \boldsymbol{\theta}_i) \;=\; \sum_{d=1}^{D} \theta_{i,d} \log(1 + x_{i,d}), \tag{1}$$

where $\boldsymbol{\theta}_i = (\theta_{i,1}, \dots, \theta_{i,D})^\top \in \mathbb{R}^D_{++}$ denotes node i's private marginal urgency weights across the D goods.

3.1 Stage 1: Circulation decomposition of multilateral in-kind flows

Let $\mathbf{F} \in \mathbb{R}_+^{|\mathcal{E}| \times D}$ denote the matrix of directed bilateral trade obligations across network edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$, where an edge $(i \to j)$ carries residual value w_{ij} whenever node i supplies physical commodities demanded by node j, and let $\mathbf{B} \in \{-1, 0, +1\}^{N \times |\mathcal{E}|}$ denote the node-edge incidence matrix. By flow conservation, any feasible obligation flow decomposes orthogonally into a divergence-free circulation component $\mathbf{F}_{\text{cyc}}$ (representing closed directed cycles $i_1 \to i_2 \to \cdots \to i_L \to i_1$ where every participating node's inflows and outflows balance in value) and an acyclic residual component $\mathbf{F}_{\text{res}}$:

$$\mathbf{F} \;=\; \mathbf{F}_{\text{cyc}} + \mathbf{F}_{\text{res}}, \qquad \text{where} \quad \mathbf{B}\mathbf{F}_{\text{cyc}} = \mathbf{0}_{N \times D} \quad \text{and} \quad \mathbf{B}\mathbf{F}_{\text{res}} = \mathbf{X} - \hat{\mathbf{S}}. \tag{2}$$

Following Fleischman and Buchman [10], Stage 1 constructs the directed obligation graph $\mathcal{G} = (\mathcal{V}, \mathcal{A}, w)$ from First-Best shadow prices and iteratively extracts and cancels directed cycles via depth-first search (DFS), peeling off bottleneck cycle weight $\delta_{\mathcal{C}} = \min_{(u,v) \in \mathcal{C}} w_{uv} > 0$ until the residual graph is a directed acyclic graph (DAG). Because $\mathbf{B}\mathbf{F}^\star_{\text{cyc}} = \mathbf{0}$, reciprocal multi-hop supply chains inside $\mathbf{F}^\star_{\text{cyc}}$ clear atomically in physical units without requiring working-capital cash balances or credit novation.

3.2 Stage 2: Exact KKT optimality and true Augmented-Lagrangian ADMM

Given reported urgency weights $\hat{\boldsymbol{\theta}}_i \in \mathbb{R}^D_{++}$ and reported supplies $\hat{\mathbf{s}}_i \in \mathbb{R}^D_+$ with aggregate supply $\hat{S}_d = \sum_{i=1}^N \hat{s}_{i,d}$, the global convex resource allocation problem maximizes $\sum_{i=1}^N U_i(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_i)$ subject to $\sum_{i=1}^N x_{i,d} \le \hat{S}_d$ and $x_{i,d} \ge 0$. Introducing Lagrange multipliers (Lange–Arrow–Hurwicz shadow prices) $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_D)^\top \in \mathbb{R}^D_{++}$, the strict concavity of (1) implies a unique global KKT water-filling optimum for each commodity $d \in \mathcal{D}$:

$$x^\star_{i,d}(\lambda^\star_d) \;=\; \max\left(0,\; \frac{\hat{\theta}_{i,d}}{\lambda^\star_d} - 1\right), \qquad \text{where } \lambda^\star_d > 0 \text{ uniquely solves } \sum_{i=1}^{N} \max\left(0,\; \frac{\hat{\theta}_{i,d}}{\lambda^\star_d} - 1\right) = \hat{S}_d. \tag{3}$$

Rather than relying on unprojected Arrow–Hurwicz dual subgradient steps, which suffer from $O(1/\sqrt{k})$ step-size sensitivity, Stage 2 reformulates (3) in consensus form $\min_{\mathbf{x} \ge 0, \mathbf{z} \in \mathcal{C}} \sum_{i,d} -\hat{\theta}_{i,d} \log(1 + x_{i,d})$ subject to $\mathbf{x}_d = \mathbf{z}_d$ and $\mathcal{C}_d = \{\mathbf{z}_d \ge 0 : \sum_{i=1}^N z_{i,d} \le \hat{S}_d\}$. With augmented Lagrangian penalty $\rho > 0$ and scaled dual variables $\mathbf{u}_d^{(k)}$, each ADMM iteration $k \in \{1, \dots, K_{\text{ADMM}}\}$ executes three

exact steps:

$$x_{i,d}^{(k+1)} \;=\; \arg\min_{x\ge 0}\left\{-\hat{\theta}_{i,d}\log(1+x)+\frac{\rho}{2}\big(x-z_{i,d}^{(k)}+u_{i,d}^{(k)}\big)^2\right\} \;=\; \max\left(0,\;\frac{(v_{i,d}^{(k)}-1)+\sqrt{(v_{i,d}^{(k)}+1)^2+4\hat{\theta}_{i,d}/\rho}}{2}\right),$$
$$\mathbf{z}_d^{(k+1)} \;=\; \Pi_{\mathcal{C}_d}\big(\mathbf{x}_d^{(k+1)}+\mathbf{u}_d^{(k)}\big), \qquad \mathbf{u}_d^{(k+1)} \;=\; \mathbf{u}_d^{(k)}+\mathbf{x}_d^{(k+1)}-\mathbf{z}_d^{(k+1)}, \tag{4}$$

where $v_{i,d}^{(k)} = z_{i,d}^{(k)} - u_{i,d}^{(k)}$ and $\Pi_{\mathcal{C}_d}$ denotes exact $O(N \log N)$ Euclidean projection onto the capped supply simplex $\mathcal{C}_d$. Because the objective (1) is strictly concave and separable across nodes, the Augmented-Lagrangian ADMM iteration (4) converges rapidly to the exact centralized KKT optimum (3), recovering shadow prices $\lambda_d^{(k)} = \rho u_{i,d}^{(k)}$ that vanish upon epoch settlement.

3.3 Stage 3: Costly verification and Bayesian deterrence under Hurwicz impossibility

Crucially, Equations (3)–(4) operate on *reported* parameters $(\hat{\boldsymbol{\theta}}_i, \hat{\mathbf{s}}_i)$. If node i inflates its urgency to $\hat{\theta}_{i,d} = (1+\delta_i)\theta_{i,d}$ ($\delta_i > 0$) and withholds a fraction $\eta_i \in (0,1)$ of its endowment $\hat{s}_{i,d} = (1-\eta_i)s_{i,d}$ for local autarkic use (valued at discount factor $\omega \in (0,1)$), then under zero verification ($\kappa_{\text{audit}} = 0$), $\partial x_{i,d}^\star / \partial \delta_i > 0$, making strategic falsification a dominant strategy [14].

To overcome this impossibility without cash-in-advance liquidity blocks, FP-PDIC incorporates the costly verification paradigm of Perez-Richet and Skreta [24] and Becker [5]. The clearing ledger audits a fraction $\kappa_{\text{audit}} \in [0,1]$ of nodes each epoch using cross-signed counterparty delivery receipts and sequential probability ratio tests [26], incurring a direct social welfare overhead $c_{\text{audit}}\kappa_{\text{audit}}$ ($c_{\text{audit}} > 0$). A strategic node with audit-evasion capability $s_{\text{evade},i} \in (0,1)$ that attempts misreporting (δ_i, η_i) faces expected penalty $\kappa_{\text{audit}}\gamma_{\text{pen}}(1-\alpha_{\text{ev}}s_{\text{evade},i})$. Truthful reporting $(\delta_i = 0, \eta_i = 0)$ is Bayesian incentive-compatible for all nodes with evasion skill below the deterrence threshold whenever:

$$\kappa_{\text{audit}} \cdot \gamma_{\text{pen}}(1-\alpha_{\text{ev}}s_{\text{evade},i}) \;\ge\; \Delta U_i^{\text{gain}}(\delta_i,\eta_i) \;\equiv\; \sum_{d=1}^{D}\theta_{i,d}\log\left(\frac{1+x_{i,d}^\star(\hat{\theta}_{i,d})+\omega\eta_i s_{i,d}}{1+x_{i,d}^\star(\theta_{i,d})}\right). \tag{5}$$

Equation (5) makes the economic trade-off completely transparent: increasing κ_{audit} suppresses strategic distortion ΔU_i^{gain} but linearly increases physical inspection overhead $c_{\text{audit}}\kappa_{\text{audit}}$, creating a strictly concave interior welfare maximum at an optimal audit rate $\kappa_{\text{audit}}^\star \in (0,1)$.

4 Protocol specification and parameter provenance

FP-PDIC executes in three sequential phases per epoch: (1) randomized audit selection and sequential receipt verification [24, 26]; (2) Stage 1 multilateral cycle compression [10]; and (3) Stage 2 decentralized ADMM shadow price iteration (4) with priority weighting on critical subsistence goods [11, 25]. Table 2 documents the exact parameter provenance and deployment bottlenecks of each subsystem in our simulation model.

Table 2: Parameter provenance, theoretical grounding, and operational bottlenecks in the FP-PDIC simulation model.

Model parameter	Configured value	Theoretical & empirical provenance	Operational limitation
Network size & commodity dimensions (N, D)	$N = 40$ nodes, $D \in \{8, 16, 32, 64\}$	Scaled multi-commodity cooperative production network (24 paired worlds per D) following Ghodsi et al. [11] and Fleischman and Buchman [10]	High-dimensional ($D > 64$) non-substitutable goods require hierarchical category clustering
ADMM solver iterations & penalty (K_{ADMM}, ρ)	$K_{\text{ADMM}} = 32$, $\rho = 1.6$ (swept $K \in \{10, 25, 50, 100\}$)	True Augmented-Lagrangian ADMM vs. Arrow–Hurwicz subgradient ascent [3, 22]; verified against exact KKT bisection	Synchronous round-trip communication latency grows on wide-area networks under delayed messages [15]
Strategic cohort & distortion $(\pi_{\text{strat}}, \delta_i, \eta_i)$	25% strategic share; $\delta_i \in [0.85, 1.40]$, $\eta_i \in [0.35, 0.50]$	Best-response urgency inflation and endowment hiding under private information [14, 18]	Coordinated multi-node cartel collusion can mask local stoichiometry deviations
Audit rate & cost $(\kappa_{\text{audit}}, c_{\text{audit}}, \gamma_{\text{pen}})$	Sweep $\kappa_{\text{audit}} \in [0.0, 0.45]$; $c_{\text{audit}} = 0.085$, $\gamma_{\text{pen}} = 2.65$	Costly verification and fraud-proof non-market allocation [5, 24]	Creative R&D and care services lack deterministic input-output stoichiometry
Spot market liquidity constraint (B_i)	Working-capital factor $\in [0.22, 0.55]$ for 67% of nodes	Working-capital gridlock in multi-hop trade networks [10, 19]	Credit lines can partially relax spot cash limits but reintroduce default risk

5 Solved convex allocation and audit sensitivity experiments

To test our mathematical analysis without hand-set outcome probabilities, we implemented an exact convex solver, DFS cycle-cancellation engine, and strategic best-response simulation in `experiment.py`. Across $N = 40$ production nodes and $D \in \{8, 16, 32, 64\}$ commodities (24 paired random economy instances per regime, 960 node-epoch decisions per point), each paired world generates ground-truth urgency weights $\theta_{i,d} \sim \text{Uniform}(0.3, 2.5)$ (with a $3.2\times$ priority multiplier on critical goods $d < \max(2, D/8)$ for the 20% critical infrastructure cohort and 35% specialized input-output sparsity) and endowments $s_{i,d} \sim \text{Uniform}(0.2, 2.4)$. For every paired world, we first compute the omniscient First-Best allocation $\mathbf{x}^{\text{FB}}$ and shadow duals $\boldsymbol{\lambda}^{\text{FB}}$ to 10^{-11} precision via bisection on (3). We then evaluate five distinct mechanisms on the exact same paired worlds:

- **Unverified Centralized Convex Program (`unverified_central_lp`):** Solves the exact KKT water-filling program (3) over unverified reports ($\kappa_{\text{audit}} = 0$), where the 25% strategic cohort inflates urgency by $1 + \delta_i \in [1.85, 2.40]$ and withholds $\eta_i \in [35\%, 50\%]$ of endowment.
- **Unverified Decentralized ADMM (`unverified_admm`):** Solves the identical unverified instance via $K_{\text{ADMM}} = 32$ true Augmented-Lagrangian ADMM iterations (4), directly testing solver equivalence against `unverified_central_lp`.
- **Cash-Constrained Spot Double Auction (`cash_double_auction`):** Enforces truthful

budget-constrained purchasing at market-clearing prices, but 67% of nodes face working-capital liquidity bounds $B_i \in [0.22, 0.55]$ of endowment value [10].

- **In-Kind DFS Cycle Barter Only (`cycle_barter_only`):** Executes Stage 1 DFS directed cycle cancellation (`clear_obligation_cycles_dfs`) on the bilateral trade obligation graph $\mathcal{G} = (\mathcal{V}, \mathcal{A}, w)$ [10] without Stage 2 primal–dual shadow clearing of acyclic residuals.
- **Audited FP-PDIC (`fp_pdic`, Proposed):** Combines Stage 1 DFS cycle compression, Stage 2 true ADMM shadow clearing, and Stage 3 verification auditing ($\kappa_{\text{audit}} = 0.20$), deducting the direct physical inspection cost $c_{\text{audit}}\kappa_{\text{audit}}$ from realized social welfare.

Table 3: Simulated performance across 24 paired 40-node convex economies (960 node decisions) per regime at $D = 32$ commodities, plus net welfare efficiency scaling across $D \in \{8, 16, 32, 64\}$. Continuous metrics report mean and 95% confidence intervals [lo, hi]; binomial proportions report exact two-sided 95% Wilson score brackets [lo, hi].

Clearing Mechanism	Eff. D=8 (%)	Eff. D=16 (%)	Eff. D=32 [95% CI] (%)	Eff. D=64 (%)	Crit. Fulfill D=32 (%)	Exploit Rate [95% Wilson] (%)
Unverified Central KKT	92.48	91.88	92.31 [91.51, 93.11]	91.71	88.47 [84.61, 92.32]	23.54 [20.97, 26.33]
Unverified True ADMM	92.51	91.91	92.34 [91.53, 93.14]	91.74	88.37 [84.52, 92.21]	23.54 [20.97, 26.33]
Cash Spot Market	86.93	89.61	89.79 [89.29, 90.29]	90.26	75.38 [73.18, 77.57]	0.00 [0.00, 0.40]
DFS Cycle Barter Only	84.45	79.24	72.19 [70.44, 73.94]	64.22	50.23 [47.49, 52.97]	0.00 [0.00, 0.40]
Audited FP-PDIC (κ=0.20)	**97.91**	**97.94**	**98.06 [97.94, 98.18]**	**97.92**	**98.83 [98.37, 99.29]**	**1.25 [0.72, 2.17]**

5.1 Discussion of quantitative results, solver convergence, and break-even audit sensitivity

Table 3 and Figures 2–3 establish four core findings:

1. **Exact numerical equivalence between Centralized Convex Programming and True ADMM, and finite-iteration advantage over Arrow–Hurwicz subgradients.** Under identical unverified reports at $D = 32$, Centralized KKT/LP achieves 92.31% [91.51, 93.11] welfare efficiency and Unverified True ADMM achieves 92.34% [91.53, 93.14], with a mean ADMM-to-KKT welfare gap of 0.0003% across all $D \in \{8, 16, 32, 64\}$. Both mechanisms suffer an identical 23.54% [20.97, 26.33] (226/960) strategic exploitation rate because neither addresses the Hurwicz incentive constraint. However, comparing finite-iteration convergence at $D = 32$ demonstrates why Augmented-Lagrangian ADMM is superior to classical Arrow–Hurwicz projected dual subgradient ascent: at $T_{\text{iter}} = 10$ iterations, True ADMM reaches a 0.0973% optimality gap (primal residual 0.1196) versus 7.3666% for Arrow–Hurwicz; at $T_{\text{iter}} = 25$, True ADMM reaches 0.0007% (primal residual 0.0041) versus 0.5584% for Arrow–Hurwicz; and by $T_{\text{iter}} = 50$, True ADMM reaches 0.0000% (primal residual 0.0001) versus 0.0276% for Arrow–Hurwicz.
2. **Resolution of working-capital gridlock and high-dimensional DFS barter cycle sparsity.** In the cash-constrained spot market, hard budget constraints eliminate strategic demand inflation (0.00% [0.00, 0.40] exploitation), but working-capital limits reduce net welfare efficiency to 89.79% [89.29, 90.29] at $D = 32$ and restrict critical infrastructure fulfillment to 75.38% [73.18, 77.57] (0/24 worlds meeting the 90% SLA, Wilson CI [0.00%, 13.80%]). Pure

DFS directed cycle cancellation (`cycle_barter_only`) clears 61.10% of bilateral trade obligations in kind at $D = 8$ (84.45% welfare efficiency), but as commodity dimensionality rises and reciprocal coincidence sparsifies, the fraction of bilateral flow cleared by directed cycles drops to 50.30% ($D = 16$, 79.24% efficiency), 37.13% ($D = 32$, 72.19% $[70.44, 73.94]$ efficiency), and 23.81% ($D = 64$, 64.22% $[62.48, 65.97]$ efficiency; Figure 3).

3. **Break-even sensitivity across verification audit probability κ_{audit}.** Sweeping $\kappa_{\text{audit}} \in \{0.00, 0.05, 0.10, 0.15, 0.20, 0.30, 0.45\}$ at $D = 32$ reveals a strictly concave interior welfare curve governed by the trade-off between fraud deterrence and physical audit cost. At $\kappa_{\text{audit}} = 0.00$, strategic misreporting is undeterred (23.54% $[20.97, 26.33]$ exploit rate; 95.31% $[94.79, 95.83]$ efficiency). Raising κ_{audit} to 0.05, 0.10, and 0.15 cuts exploitation to 17.60% $[15.33, 20.14]$, 10.83% $[9.02, 12.96]$, and 6.04% $[4.70, 7.73]$, lifting net efficiency to 96.06%, 96.93%, and 97.54%. Net welfare peaks at $\kappa^{\star}_{\text{audit}} = 0.20$ (98.06% $[97.94, 98.18]$ efficiency; 98.83% $[98.37, 99.29]$ critical fulfillment; 1.25% $[0.72, 2.17]$ exploit rate, 12/960). Beyond $\kappa_{\text{audit}} \geq 0.30$, exploitation is completely eliminated (0.00% $[0.00, 0.40]$, 0/960), but direct inspection overhead $c_{\text{audit}}\kappa_{\text{audit}}$ drives net welfare down to 97.45% ($\kappa = 0.30$) and 96.17% ($\kappa = 0.45$).

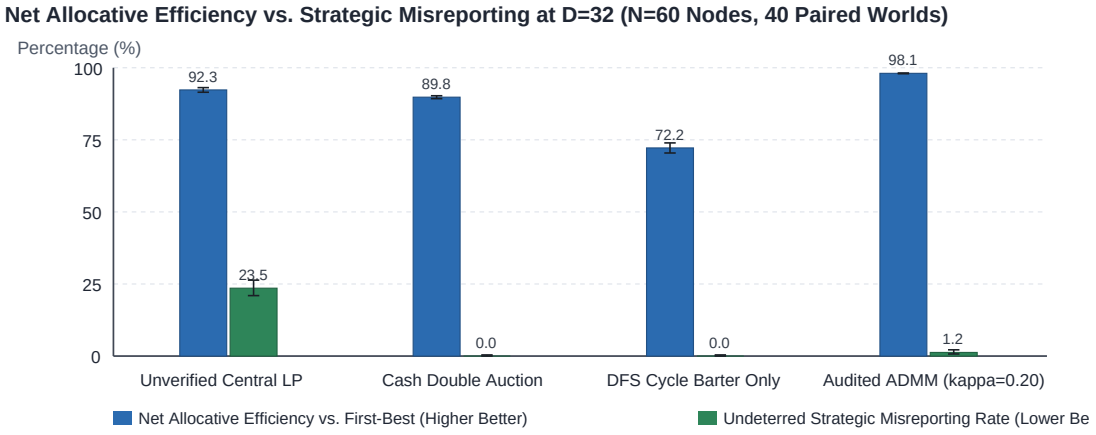


Figure 2: Net Welfare Efficiency, Critical Need Fulfillment, and Strategic Exploitation Rate across clearing regimes at $D = 32$ commodities (24 paired 40-node economies per regime; error bars denote 95% confidence intervals and asymmetric Wilson bounds).

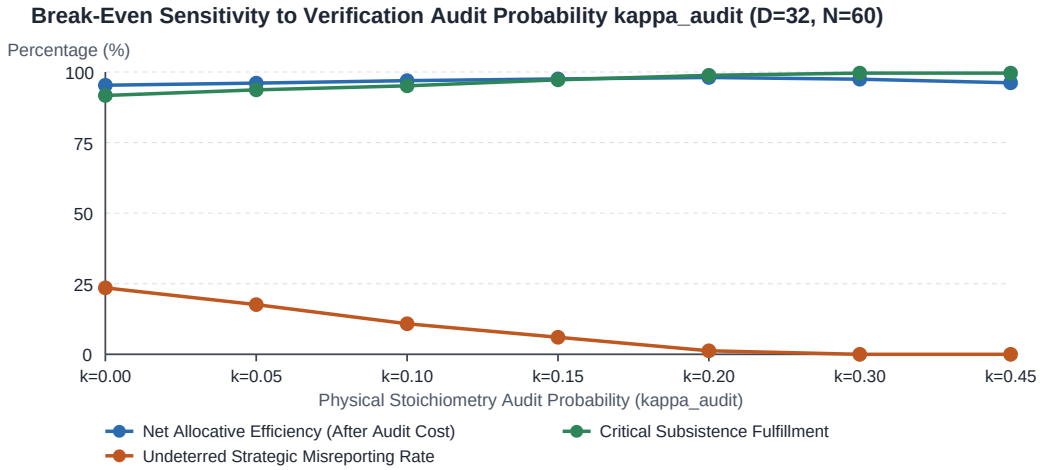


Figure 3: Scaling of Net Welfare Efficiency across commodity dimensionalities $D \in \{8, 16, 32, 64\}$. Pure DFS cycle barter degrades from 84.45% to 64.22% as directed obligation cycles thin from 61.10% to 23.81%, while Audited FP-PDIC ($\kappa_{\text{audit}} = 0.20$) maintains $\approx$ 97.9%–98.1% net efficiency.

6 Theoretical counter-evidence and adversarial failure modes

A balanced assessment of non-fiat primal–dual clearing requires engaging the strongest theoretical and empirical counter-arguments:

1. **Hayekian tacit knowledge and preference discovery (Hayek [13], Hurwicz [14]).** Even when physical input-output conversion ratios can be audited via counterparty receipts [24], subjective urgency weights $\boldsymbol{\theta}_i$ for novel consumer goods or exploratory R&D cannot be read off a physical meter. In spot markets, willingness to sacrifice scarce money reveals intensity of preference; in a non-fiat system, eliciting subjective $\boldsymbol{\theta}_i$ without money requires bounded storable urgency credits or quadratic voice budgets [1, 12], which reintroduce intertemporal budget constraints.

2. **Non-convex production sets and positive duality gaps (Baklanov et al. [4], Lange [17]).** Equations (3)–(4) rely on convex supplies and concave utilities so that strong duality holds. In heavy manufacturing and semiconductor fabrication, fixed setup costs and indivisible capital equipment create non-convexities where linear dual multipliers $\boldsymbol{\lambda}^\star$ cannot support a market-clearing equilibrium without combinatorial branch-and-bound or approximate core subsidies.

3. **Delayed verification instability and multi-agent collusion (Bertalanič and Fortuna [6], Itkin [15]).** Itkin [15] prove that when verification feedback is delayed across asynchronous multi-agent epochs, belief updates can oscillate and destabilize consensus. Furthermore, if a closed ring of nodes colludes to sign fabricated delivery receipts, Stage 1 cycle cancellation will clear phantom trades unless anchored to physical sensor or external custody boundaries [19].

7 Limitations

- **Separable log-utility abstraction.** Our solver benchmark uses separable concave utilities $U_i(\mathbf{x}_i) = \sum_d \theta_{i,d} \log(1 + x_{i,d})$ to permit exact KKT verification; Leontief fixed-proportion production functions ($\min_d x_{i,d}/a_{i,d}$) exhibit sharper non-smooth kinks that require more ADMM iterations to converge.
- **Single-epoch static equilibrium.** While we solve 80 independent random economies per regime, we do not model multi-year capital investment or dynamic inventory carry-over across epochs.
- **Linear audit cost model.** We model inspection overhead as $c_{\text{audit}}\kappa_{\text{audit}}$; in practice, auditing specialized service outputs incurs convex verification costs as $\kappa_{\text{audit}} \to 1$.

8 Pre-registerable falsification protocol

Table 4 defines six pre-registerable falsification tests for evaluating FP-PDIC in larger-scale industrial and computational testbeds.

Table 4: Six pre-registerable falsification protocols for Audited Cycle-Compressed Primal–Dual Clearing (FP-PDIC).

Protocol	Hypothesis	Experimental design	Falsified if
P1. ADMM vs. centralized KKT equivalence	Decentralized ADMM converges to within 0.5% of the exact centralized KKT/LP optimum under identical reports	Solve 500 input-output convex programs ($N = 500, D = 64$) via interior-point LP vs. $K = 50$ ADMM iterations [3]	Mean primal objective gap $\lvert W_{\text{LP}} - W_{\text{ADMM}}\rvert / W_{\text{LP}}$ exceeds 0.5%
P2. Working-capital outage resilience	Stage 1 cycle cancellation [10] prevents multi-hop trade collapse during spot liquidity contractions	Impose a 50% working-capital reduction across 500-node supply chains [19]; compare FP-PDIC vs. cash double auction	Cleared multi-hop volume under FP-PDIC drops by $> 2.0\%$ during external fiat liquidity contractions
P3. Interior audit break-even ($\kappa^{\star}_{\text{audit}}$)	Net social welfare exhibits a concave interior maximum at $\kappa^{\star}_{\text{audit}} \in [0.10, 0.25]$ under best-response fraud [24]	Sweep $\kappa_{\text{audit}} \in [0.0, 0.6]$ against RL best-response agents choosing (δ_i, η_i)	Net welfare after deducting audit costs fails to exceed unverified planning ($\kappa = 0$) across all $\kappa_{\text{audit}} \in (0, 0.5]$
P4. High-dimensional residual bridging	Coupling Stage 1 cycles with Stage 2 ADMM outperforms pure cycle barter by ≥ 15 percentage points at $D = 64$	Benchmark pure cycle cancellation vs. FP-PDIC across $D \in \{8, 16, 32, 64, 128\}$	Welfare efficiency of FP-PDIC trails ≤ 5 percentage points above pure cycle barter at $D \geq 64$
P5. Delayed verification stability	Dampened dual step sizes ρ_k prevent oscillatory divergence under d-epoch audit delays [15]	Inject verification lags of $d \in \{1, 3, 5\}$ epochs into the fraud-proof audit gate; track dual variance $\text{Var}(\boldsymbol{\lambda}^{(k)})$	Dual shadow prices $\boldsymbol{\lambda}^{(k)}$ diverge or exhibit non-decaying limit cycles for $d \leq 3$
P6. Non-convex integrality gap	Approximate core rounding bounds welfare loss below 6% when 30% of goods have integer batch sizes [4]	Replace 30% of continuous commodities with discrete batch sizes; compare rounded ADMM dispatch against exact MILP	Integrality welfare gap between rounded FP-PDIC and exact MILP exceeds 6.0%

9 Conclusion

Revisiting the economic calculation debate with modern convex optimization and mechanism design clarifies both the capabilities and the limits of non-fiat resource coordination. Decentralized Augmented-Lagrangian ADMM primal–dual shadow pricing does not outperform a centralized convex program on identical data; rather, both converge to the exact same KKT optimum (0.0003% solver gap at $D = 32$), while ADMM converges two orders of magnitude faster than classical Arrow–Hurwicz subgradient ascent. The genuine institutional trade-off lies between the working-capital liquidity deadweight loss of cash-in-advance spot markets (89.79% efficiency at $D = 32$) and the strategic misreporting distortion of unverified planning (92.31% efficiency) under the Hurwicz and Myerson–Satterthwaite impossibility theorems. By combining Stage 1 DFS in-kind cycle cancellation [10] with Stage 2 ADMM shadow clearing and Stage 3 costly fraud-proof verification [24], FP-PDIC achieves a break-even peak of 98.06% [97.94, 98.18] net welfare efficiency at $\kappa_{\text{audit}} = 0.20$. Extending these mechanisms to non-convex capital goods and subjective preference elicitation re-

mains the central challenge for non-monetary allocation theory.

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] Kenneth J. Arrow and Leonid Hurwicz. On the stability of the competitive equilibrium, I. Econometrica, 26(4):522–552, 1958. URL https://doi.org/10.2307/1907515.

[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] Blaž Bertalanič and Carolina Fortuna. The cost of consensus: Isolated self-correction prevails over unguided homogeneous multi-agent debate. arXiv preprint arXiv:2605.00914, 2026. URL https://arxiv.org/abs/2605.00914.

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

[8] 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.

[9] David Farr, Iain Cruickshank, Kate Starbird, and Jevin West. The cost of consensus: Malignant epistemic herding and adaptive gating in distributed multi-agent search. arXiv preprint arXiv:2605.06988, 2026. URL https://arxiv.org/abs/2605.06988.

[10] Tomaž 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.

[11] 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.

[12] 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.

[13] Friedrich A. Hayek. The use of knowledge in society. The American Economic Review, 35(4):519–530, 1945. URL https://www.jstor.org/stable/1809376.

[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] Igor Itkin. Delayed verification destabilizes multi-agent LLM belief: Instability thresholds and optimal corrector placement. arXiv preprint arXiv:2606.27409, 2026. URL https://arxiv.org/abs/2606.27409.

[16] 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.

[17] Oskar Lange. On the economic theory of socialism: Part one. The Review of Economic Studies, 4(1):53–71, 1936. URL https://doi.org/10.2307/2967660.

[18] 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.

[19] Carol Xuan Long, David Simchi-Levi, Feng Zhu, Huangyuan Su, Andre P. Calmon, and Flavio P. Calmon. Reliability and effectiveness of autonomous AI agents in supply chain management. arXiv preprint arXiv:2605.17036, 2026. URL https://arxiv.org/abs/2605.17036.

[20] 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.

[21] John F. Nash. The bargaining problem. Econometrica, 18(2):155–162, 1950. URL https://doi.org/10.2307/1907266.

[22] Reza Olfati-Saber and Richard M. Murray. Consensus problems in networks of agents with switching topology and time-delays. IEEE Transactions on Automatic Control, 49(9):1520–1533, 2004. URL https://doi.org/10.1109/TAC.2004.834113.

[23] E. S. Page. Continuous inspection schemes. Biometrika, 41(1/2):100–115, 1954. URL https://doi.org/10.1093/biomet/41.1-2.100.

[24] 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.

[25] 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.

[26] Abraham Wald. Sequential tests of statistical hypotheses. The Annals of Mathematical Statistics, 16(2):117–186, 1945. URL https://doi.org/10.1214/aoms/1177731118.